

Practical Tikhonov Regularized Estimators in Reproducing Kernel Hilbert Spaces for Statistical Inverse Problems

Robert Hable

Abstract

Regularized kernel methods such as support vector machines (SVM) and support vector regression (SVR) constitute a broad and flexible class of methods which are theoretically well investigated and commonly used in nonparametric classification and regression problems. As these methods are based on a Tikhonov regularization which is also common in inverse problems, this article investigates the use of regularized kernel methods for inverse problems in a unifying way. Regularized kernel methods are based on the use of reproducing kernel Hilbert spaces (RKHS) which lead to very good computational properties. It is shown that similar properties remain true in solving statistical inverse problems and that standard software implementations developed for ordinary regression problems can still be used for inverse regression problems.

Consistency of these methods and a rate of convergence for the risk is shown under quite weak assumptions and rates of convergence for the estimator are shown under somehow stronger assumptions. The applicability of these methods is demonstrated in a simulation.

1 Introduction

One of the most important statistical inverse problems is the inverse regression problem in which one observes i.i.d. data $(z_1, y_1), \dots, (z_n, y_n)$ from the model

$$Y = (Af_0)(Z) + s(Z)\varepsilon \quad (1)$$

in which A is a (known) linear operator between suitable function spaces, $Af_0 : \mathcal{Z} \rightarrow \mathbb{R}$ is the (unknown) regression function, and $s : \mathcal{Z} \rightarrow \mathbb{R}$ is an (unknown) scale function. The goal is to estimate the primary function $f_0 : \mathcal{X} \rightarrow \mathbb{R}$. If A has a bounded inverse A^{-1} , then f_0 can simply be estimated by $A^{-1}\hat{g}_n$ where \hat{g}_n is an ordinary estimate of the regression function $g = Af_0$. However, in a typical inverse regression problem, A does not have a bounded inverse so that one is faced with an ill-posed problem which has to be dealt

with in different and much more complicated ways. See, e.g., Cavalier (2011) for an overview. Two of the most common types of estimators in inverse regression problems are spectral cut-off estimators and Tikhonov estimators. In case of the spectral cut-off estimator, it is assumed that A is an injective compact operator between L_2 -spaces $L_2(\mu)$ and $L_2(P_Z)$ so that a singular value decomposition exists. That is, there are a complete orthonormal system $(v_j)_{j \in \mathbb{N}}$ of $L_2(\mu)$, an orthonormal system $(u_j)_{j \in \mathbb{N}}$ in $L_2(P_Z)$, and singular values $(\sigma_j)_{j \in \mathbb{N}} \subset (0, \infty)$ such that $Av_j = \sigma_j u_j$ and $A^* u_j = \sigma_j v_j$ where A^* denotes the adjoint operator of A . Then, the spectral cut-off estimator is given by

$$\hat{f}_n(x) = \sum_{j=1}^J \frac{\hat{b}_{j,n}}{\sigma_j} v_j(x) \quad \text{with} \quad \hat{b}_{j,n} = \frac{1}{n} \sum_{i=1}^n u_j(z_i) y_i$$

where the truncation parameter $J = J_n$ acts as a regularization parameter. This estimator is investigated in many articles on inverse regression problems, e.g., in van Rooij and Ruymgaart (1996), Mair and Ruymgaart (1996), Bauer and Munk (2007), Bissantz and Holzmann (2008), and Bissantz and Birke (2009). A disadvantage of this estimator is that one needs to know the singular value decomposition of A in order to compute the estimator (this is also a considerable limitation for general software implementations). Furthermore, one usually has to know the distribution P_Z of the covariate Z . Most articles on spectral cut-off estimators assume that Z is uniformly distributed on $[0, 1]$ or use equidistant design points.

Methods based on Tikhonov regularizations are common in non-stochastic as well as in statistical settings of inverse problems. In inverse regression problems, the estimator based on Tikhonov regularization is the minimizer

$$\arg \min_{f \in H} \left(\frac{1}{n} \sum_{i=1}^n (y_i - (Af)(z_i))^2 + \lambda \|f\|_H^2 \right) \quad (2)$$

where H is a Hilbert space of functions $f : \mathcal{X} \rightarrow \mathbb{R}$. Such an estimator has been considered, e.g., in O'Sullivan (1986), Mathé and Pereverzev (2001), Bissantz et al. (2007), and Cavalier (2008). Most articles on Tikhonov regularization in statistical inverse problems focus on rates of convergence, but simulations or applications on real data sets are only rarely done. One reason might be that calculating the estimator is, in general, not an easy task and suitable software implementations (e.g., as R-packages) are still widely missing. The situation gets better if reproducing kernel Hilbert spaces (RKHS) are chosen for H in (2). These Hilbert spaces have excellent properties from a computational point of view and, therefore, recently attract much attention in statistics, machine learning, and approximation theory. Tikhonov estimators (2) in an RKHS are already used in the early work Wahba (1977) and Wahba (1980); there, a special case of model (1) is considered in which the error is homoscedastic (i.e., $s \equiv 1$), $\mathcal{X} = \mathcal{Z} = [0, 1]$, the data $z_i, i \in \{1, \dots, n\}$,

are equidistant design points, and A is a Fredholm integral operator of the first kind, that is,

$$Af(z) = \int K(x, z)f(x) \lambda(dx) \quad \forall z \in [0, 1], \quad f \in H.$$

A similar setting is also considered in Nychka and Cox (1989) which shows rates of convergence for Tikhonov estimators (2) in an RKHS. In ordinary nonparametric classification and regression problems, lots of research has been done on methods based on RKHS such as support vector machines (SVM) and support vector regression (SVR) during the last decade. These methods belong to a broad class of methods called regularized kernel methods. In an ordinary (heteroscedastic) regression problem

$$Y = f_0(X) + s(X)\varepsilon, \quad (3)$$

the estimate of a regularized kernel method is the minimizer

$$\arg \min_{f \in H} \left(\frac{1}{n} \sum_{i=1}^n L(x_i, y_i, f(x_i)) + \lambda \|f\|_H^2 \right) \quad (4)$$

where H is an RKHS and L is a suitable loss function; see, e.g., Vapnik (1998), Schölkopf and Smola (2002), and Steinwart and Christmann (2008). In view of (2), regularized kernel methods can also be defined for inverse regression problems in an obvious way by

$$\arg \min_{f \in H} \left(\frac{1}{n} \sum_{i=1}^n L(z_i, y_i, (Af)(z_i)) + \lambda \|f\|_H^2 \right). \quad (5)$$

The goal of the present article is to investigate these methods (which considerably generalize the setting of the early work in Wahba (1977), Wahba (1980), and Nychka and Cox (1989)) in a unifying way in the light of the current state of research on regularized kernel methods for ordinary regression problems (3). The results are not restricted to $\mathcal{X} = \mathcal{Z} = [0, 1]$ but \mathcal{X} may be any compact subset of \mathcal{R}^d for any dimension $d \in \mathbb{N}$ and \mathcal{Z} may be any Polish space. Furthermore, we also consider heteroscedastic errors as homoscedasticity is less frequent in real data sets. Allowing for different loss functions, on the one hand, extends the applicability of the method from mean regression to tasks such as median regression, quantile regression, and classification. On the other hand, it is well known that the least-squares loss typically induces bad robustness properties and that regularized kernel methods for Lipschitz-continuous loss functions (such as the absolute deviation loss) have very good robustness properties; see Christmann and Steinwart (2007), Christmann et al. (2009), and Hable and Christmann (2011). Formally, using the least-squares loss simply was a computational need; due

to the nice structure of the least-squares loss, calculating (2) is equivalent to solving the equality

$$A^*Af + n\lambda f = A^*y, \quad f \in H, \quad (6)$$

in H and, for compact operators A with singular system $(\sigma_j; v_j, u_j)$, the solution is given by

$$\hat{f}_n = \sum_{j=1}^{\infty} \frac{\sigma_j}{\sigma_j^2 + n\lambda} \langle y, u_n \rangle v_j;$$

see, e.g., (Engl et al., 1996, p. 117). If H is an RKHS, solving (6) is much easier but still involves calculating the inverse of an $n \times n$ -matrix; see (Wahba, 1977, p. 654). However, nowadays, enormous efforts have been made in order to develop powerful software implementations for calculating regularized kernel methods (4) such as SVM and SVR for various loss functions. It turns out (Theorem 2.3) that these implementations can still be used in order to calculate (5), i.e., regularized kernel methods for inverse problems. One only has to calculate a certain “pseudo” kernel matrix M and proceed with standard software implementations as if M was the kernel matrix K from an ordinary regularized kernel method; see Section 2 and 4 for details.

In a non-stochastic setting, an RKHS in (5) has also been considered in Krebs et al. (2009) for the least-squares loss and in Krebs (2011) for the ε -insensitive loss, a common loss function in machine learning. Also in a non-stochastic setting, Eggermont et al. (2012) consider an RKHS in the Tikhonov method (with the least-squares loss) but in a quite different way: there, the codomain of A is the subset of an RKHS while we use an RKHS as the domain of A .

The rest of the article is organized as follows: Section 2 contains notations, assumptions, and the general definition of regularized kernel methods for inverse problems. Furthermore, it is shown that the estimators uniquely exist (Theorem 2.4) and that an analogue of the empirical representer theorem holds which enables to use standard software implementations developed for ordinary regularized kernel methods (Theorem 2.3). In Section 3, consistency in the H -norm (which is stronger than the supremum-norm) and a rate of convergence for the risk is shown under quite weak assumptions (Theorem 3.1 and Theorem 3.2). A rate of convergence of the estimator in the H -norm is shown under somehow stronger assumptions (Theorem 3.3). Section 4 contains a simulation in a standard example, namely the heat equation. All proofs are deferred to the appendix; Subsection 6.1 in the appendix also contains a number of additional results which are needed in the proofs of the main results and are interesting on its own.

2 Regularized Kernel Methods in Inverse Problems

Throughout the whole article, we deal with the general setting summarized in the following assumption:

Assumption 2.1 *Let P be a probability measure on $(\mathcal{Z} \times \mathcal{Y}, \mathfrak{B}_{\mathcal{Z} \times \mathcal{Y}})$ where \mathcal{Z} is a Polish space, $\mathcal{Y} \subset \mathbb{R}$ is closed and $\mathfrak{B}_{\mathcal{Z} \times \mathcal{Y}}$ is the Borel- σ -algebra on $\mathcal{Z} \times \mathcal{Y}$. The marginal distribution of P on \mathcal{Z} is denoted by $P_{\mathcal{Z}}$; the corresponding conditional distribution on \mathcal{Y} given z is denoted by $P(\cdot|z)$. That is,*

$$\int g dP = \iint g(z, y) P(dy|z) P_{\mathcal{Z}}(dz) \quad \forall g \in L_1(P).$$

Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and let k be a kernel on \mathcal{X} which is extendable to an m -times differentiable kernel on \mathbb{R}^d where $m > \frac{d}{2}$. The RKHS of k is denoted by H . The operator

$$A : H \rightarrow \mathcal{C}_b(\mathcal{Z}) \quad \text{is continuous and linear} \quad (7)$$

where $\mathcal{C}_b(\mathcal{Z})$ denotes the Banach space of all bounded, continuous functions $g : \mathcal{Z} \rightarrow \mathbb{R}$ with supremum-norm $\|\cdot\|_{\infty}$.

The function $L : \mathcal{Z} \times \mathcal{Y} \times \mathbb{R} \rightarrow [0, \infty)$ is a continuous loss function; that is, the function $t \mapsto L(z, y, t)$ is convex for every fixed $z \in \mathcal{Z}, y \in \mathcal{Y}$.

In order to prove consistency and rates of convergence, Assumption (7) will be replaced by the stronger assumption that

$$A : \mathcal{C}_b(\mathcal{X}) \rightarrow \mathcal{C}_b(\mathcal{Z}) \quad \text{is continuous and linear} \quad (8)$$

later on. However, it is always made explicit whenever (8) is assumed instead of (7). Note that Assumption (8) indeed implies (7) because, according to (Steinwart and Christmann, 2008, Lemma 4.23),

$$\|f\|_{\infty} \leq \|k\|_{\infty} \|f\|_H \quad \forall f \in H. \quad (9)$$

Articles on inverse regression problems typically assume compactness of the operator A and it seems as if we had no such assumption here. However, A enjoys a compactness property which comes for free in this setting. As shown in Prop. 6.1, it follows from (8) that $A : H \rightarrow \mathcal{C}_b(\mathcal{Z})$ is a compact operator, but this compactness is a quite weak property. The reason for this is that weak convergence in the RKHS of a bounded continuous kernel is a relatively strong kind of convergence. In particular, weak convergence in H implies pointwise convergence, which is an easy consequence of the so-called reproducing property

$$\langle f, \Phi(x) \rangle_H = f(x) \quad \forall x \in \mathcal{X}, f \in H \quad (10)$$

where Φ denotes the canonical feature map of H , i.e., $\Phi(x) = k(x, \cdot)$ for every $x \in \mathcal{X}$.

In most parts of the article, we will also impose the following assumption on the loss function L :

Assumption 2.2 *Assume that L is continuous and that there are $b \in \mathcal{L}_1(P)$, $c_L \in (0, \infty)$, and $\beta_L \in [1, 2]$ such that, for every $z \in \mathcal{Z}$, $y \in \mathcal{Y}$, and $t \in \mathbb{R}$,*

$$L(z, y, t) \leq b(z, y) + c_L |t|^{\beta_L} . \quad (11)$$

In addition, assume that there are $p \in [0, 1]$, $b'_0 \in \mathcal{L}_2(P_{\mathcal{Y}})$ with $b'_0 \geq 0$, and $b'_1 \in [0, \infty)$ such that, for every $z \in \mathcal{Z}$, $y \in \mathcal{Y}$, $a \in (0, \infty)$, and $t_1, t_2 \in [-a, a]$,

$$|L(z, y, t_1) - L(z, y, t_2)| \leq (b'_0(y) + b'_1 a^p) \cdot |t_1 - t_2| . \quad (12)$$

This assumption looks quite special but, indeed, covers all of the commonly used loss functions: least squares, hinge, truncated least squares, and logistic for classification; least squares, absolute distance, pinball, epsilon-insensitive, Huber, and logistic for regression.

For $D_n = ((z_1, y_1), \dots, (z_n, y_n)) \in (\mathcal{Z} \times \mathcal{Y})^n$ and $\lambda > 0$, define

$$f_{A, D_n, \lambda} = \arg \min_{f \in H} \left(\frac{1}{n} \sum_{i=1}^n L(z_i, y_i, (Af)(z_i)) + \lambda \|f\|_H^2 \right) \quad (13)$$

and the regularized empirical risk

$$\mathcal{R}_{A, D, \lambda}(f) = \frac{1}{n} \sum_{i=1}^n L(z_i, y_i, (Af)(z_i)) + \lambda \|f\|_H^2 \quad \forall f \in H .$$

That is, $f_{A, D_n, \lambda} = \arg \min_{f \in H} \mathcal{R}_{A, D, \lambda}(f)$. The following theorem is the analogue to the well-known representer theorem in case of ordinary regularized kernel methods (i.e. $A = \text{id}$); see, e.g., (Steinwart and Christmann, 2008, Theorem 5.5). By use of this theorem, the optimization problem (13) in the infinite-dimensional function space H can be reduced to a convex optimization problem in \mathbb{R}^n . In similar but non-stochastic inverse problems, corresponding results have already been obtained for special loss functions in (Krebs et al., 2009, Lemma 3.1 and Theorem 3.2) and (Krebs, 2011, Lemma 3.1).

Theorem 2.3 (Empirical Representer Theorem)

Let Assumption 2.1 be fulfilled. For every $D_n = ((z_1, y_1), \dots, (z_n, y_n)) \in$

$(\mathcal{Z} \times \mathcal{Y})^n$ and $\lambda > 0$, the function $f_{A,D_n,\lambda}$ defined by (13) uniquely exists. There are $\alpha_1, \dots, \alpha_n \in \mathbb{R}$ such that

$$f_{A,D_n,\lambda}(\cdot) = \sum_{i=1}^n \alpha_i \cdot (A\Phi(\cdot))(z_i). \quad (14)$$

The matrix $M \in \mathbb{R}^{n \times n}$ defined by

$$M_{i,j} = \left(A \left((A\Phi(\cdot))(z_i) \right) \right) (z_j) \quad (15)$$

is symmetric and positive semi-definite. For every $\alpha = (\alpha_1, \dots, \alpha_n)^\top \in \mathbb{R}^n$,

$$\mathcal{R}_{A,D,\lambda} \left(\sum_{i=1}^n \alpha_i \cdot (A\Phi(\cdot))(z_i) \right) = \frac{1}{n} \sum_{i=1}^n L(z_i, y_i, \alpha^\top M e_i) + \lambda \alpha^\top M \alpha \quad (16)$$

where e_i denotes the i -th vector in the standard basis of \mathbb{R}^n .

Theorem 2.3 is of great practical importance because it says that estimates can be calculated essentially by finding a minimizer of

$$\alpha \mapsto \frac{1}{n} \sum_{i=1}^n L(z_i, y_i, \alpha^\top M e_i) + \lambda \alpha^\top M \alpha$$

in \mathbb{R}^n where M is a symmetric and positive semi-definite matrix. This is extremely comfortable because calculating ordinary regularized kernel methods leads to an optimization problem with exactly the same structure. Accordingly, developing new software for inverse problems is unnecessary because, after calculating the matrix M , one can use standard software for regularized kernel methods such as the R-package “kernlab” (Karatzoglou et al., 2004). In order to calculate M , one only needs to write an R-function which takes a function f as an argument and returns the function Af .

Almost all articles on inverse regression problems assume the homoscedastic regression model

$$Y = (Af_{A,P})(Z) + \varepsilon. \quad (17)$$

Instead of only considering such a specific model, we use a suitable loss function L and consider the risk

$$\mathcal{R}_{A,P}(f) := \int L(z, y, (Af)(z)) P(d(z, y))$$

for functions $f : \mathcal{X} \rightarrow \mathbb{R}$ (in the domain of A). Then, the goal is to estimate a minimizer $f_{A,P}$ of this risk. In this way, it is, e.g., possible to investigate heteroscedastic inverse regression problems such as

$$Y = (Af_{A,P})(Z) + s(Z)\varepsilon, \quad (18)$$

where s is an unknown scale function or, even more general,

$$Y = (Af_{A,P})(Z) + \varepsilon_Z,$$

where $(\omega, z) \mapsto \varepsilon_z(\omega)$ is a Markov kernel. In most parts of the article, we do not make any specific assumption on the regression model but consider minimizing risks. As a theoretical tool, we also need the regularized risk

$$\mathcal{R}_{A,P,\lambda}(f) = \mathcal{R}_{A,P}(f) + \lambda \|f\|_H^2 = \int L(z, y, (Af)(z)) P(d(z, y)) + \lambda \|f\|_H^2$$

for $f \in H$ and $\lambda \in (0, \infty)$. This regularized risk is the theoretical counterpart of the regularized empirical risk $\mathcal{R}_{A,D,\lambda}$. The following theorem presents a simple condition under which a unique minimizer of the regularized risk exists. By choosing the empirical measure for P , the theorem also guarantees existence of the estimates $f_{A,D,\lambda}$.

Theorem 2.4 *Let Assumption 2.1 be fulfilled and*

$$\int L(z, y, 0) P(d(z, y)) < \infty. \quad (19)$$

Then, for every $\lambda > 0$, there is a unique minimizer $f_{A,P,\lambda}$ of $f \mapsto \mathcal{R}_{A,P,\lambda}(f)$ in H .

3 Consistency and Rate of Convergence

Define i.i.d. random variables

$$(X_i, Y_i) \sim P, \quad i \in \mathbb{N}.$$

Then, the data set D_n is a realization of the random vector

$$\mathbf{D}_n = ((X_1, Y_1), \dots, (X_n, Y_n)).$$

The following theorem guarantees consistency of regularized kernel methods for inverse regression problems under quite weak assumptions. In particular, it also covers the multivariate case as \mathcal{X} may be any compact subset of \mathcal{R}^d , it does not require homoscedasticity or any signal plus noise assumption, and we do not assume injectivity of the operator A or any properties of a singular system of A .

Theorem 3.1 *Let Assumptions 2.1 and 2.2 be fulfilled, and let A fulfill (8). Assume that*

$$\exists f^* \in H \text{ s.t. } \mathcal{R}_{A,P}(f^*) = \inf_{f \in H} \mathcal{R}_{A,P}(f) := \mathcal{R}_{A,P}^*. \quad (20)$$

Then, there is an $f_{A,P} \in H$ such that

$$\mathcal{R}_{A,P}(f_{A,P}) = \inf_{f \in H} \mathcal{R}_{A,P}(f) \quad (21)$$

and, for every sequence $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ such that $\lim_{n \rightarrow \infty} \lambda_n^{1+p/2} \sqrt{n} = \infty$,

$$\|f_{A,\mathbf{D}_n,\lambda_n} - f_{A,P}\|_H \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.} \quad (22)$$

Note that (22) also implies

$$\sup_{x \in \mathcal{X}} |f_{A,\mathbf{D}_n,\lambda_n}(x) - f_{A,P}(x)| \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.}$$

As already mentioned above, Theorem 3.1 does not require any signal plus noise assumption. Instead, it is only assumed that H contains a minimizer of the risk. In order to make this assumption more explicit in case of a signal plus noise assumption, it is exemplified for special choices of the loss function L and the RKHS H . Consider the heteroscedastic model

$$Y = Af_0(Z) + s(Z)\varepsilon \quad (23)$$

where s is an unknown scale function and $\mathbb{E}\varepsilon = 0$ in case of the least-squares loss or $\text{median}(\varepsilon) = 0$ in case of the absolute deviation loss. (In the latter case, it is additionally assumed that 0 is the unique median of ε ; error distributions which violate this assumption are extremely unusual.) Then, Assumption (20) is fulfilled if $f_0 \in H$. As we have not assumed injectivity of A so far, model (23) is not necessarily unique. It is possible that $f_0 \neq f_{A,P}$ but it follows from (21) that $Af_{A,P} = Af_0$ P_Z -a.s. so that model (23) can be rewritten as

$$Y = Af_{A,P}(Z) + s(Z)\varepsilon; \quad (24)$$

Obviously, distinguishing between models (23) and (24) is impossible. In order to prove rates of convergence for $f_{A,\mathbf{D}_n,\lambda_n} - f_{A,P}$ below, we will also assume that A is injective. Under this standard assumption, the model is unique and $f_0 = f_{A,P}$.

In a parametric setting, it is typically assumed that f_0 is linear or a polynomial. This assumption easily implies (20) if k is the linear kernel or a suitable polynomial kernel. In a nonparametric setting, the most common kernel is the Gaussian RBF kernel. However, assuming that a function f_0 lies in the corresponding RKHS H is a rather strong and inaccessible assumption which can hardly be made explicit for a practitioner – even though this RKHS is dense in $\mathcal{C}(\mathcal{X})$. Therefore, it seems advisable to use slightly different kernels in the nonparametric setting, namely Wendland kernels. These are radial kernels of the form

$$k_{d,m}(x, x') = \phi_{d,m}(\|x - x'\|) \quad \text{with} \quad \phi_{d,m}(r) = \begin{cases} p_{d,m}(r), & 0 \leq r \leq 1 \\ 0, & r > 1 \end{cases}$$

Dimension	Wendland polynomial
$d = 1$	$p_{1,1}(r) = (1 - r)^3(3r + 1)$
$d = 2$	$p_{2,2}(r) = (1 - r)^6(35r^2 + 18r + 3)$
$d = 3$	$p_{3,2}(r) = (1 - r)^6(35r^2 + 18r + 3)$
$d = 4$	$p_{4,3}(r) = (1 - r)^9(693r^3 + 477r^2 + 135r + 15)$
$d = 5$	$p_{5,3}(r) = (1 - r)^9(693r^3 + 477r^2 + 135r + 15)$

Table 1: Suitable Wendland polynomials $p_{d,m}$ for different dimensions d ; that is, $m = (d + 1)/2$ if d is odd, $m = d/2 + 1$ if d is even.

where $p_{d,m}$ is a certain polynomial of degree $\lfloor d/2 \rfloor + 3m + 1$. The polynomial is of minimal degree such that $k_{d,m}$ is m -times continuously differentiable; see (Wendland, 2005, Theorem 9.12 and 9.13). Though the shape of these kernels is very similar to that of the Gaussian RBF kernel, Wendland kernels have two advantages: First, they are compactly supported and therefore lead to sparse kernel matrices. Second, there is a simple condition on f_0 which guaranties that f_0 is contained in the corresponding RKHS $H_{d,m}$, i.e., that (20) is fulfilled. If the dimension d is odd, choose $m = (d+1)/2$ and $\gamma = d+1$. If d is even, choose $m = d/2 + 1$ and $\gamma = d + 2$. Then, f_0 is in $H_{d,m}$ if it is the restriction of a γ -times continuously differentiable function on \mathbb{R}^d . This is a consequence of the fact that the RKHS of $k_{d,m}$ is the Sobolev space $H^{d/2+m+1/2}(\mathbb{R}^d)$, that $\gamma \geq d/2 + m + 1/2$, and that \mathcal{X} is bounded; see (Wendland, 2005, Theorem 10.35 and § 10.7). For the convenience of the reader, Table 1 contains the relevant Wendland polynomials $p_{d,m}$ for dimensions $d \leq 5$; these are calculated from (Wendland, 2005, Cor. 9.15). Polynomials $p_{d,m}$ for higher dimensions can recursively be obtained from (Wendland, 2005, Theorem 9.12).

For the risk of the estimator, we obtain a rate of convergence – even under the quite weak assumptions of Theorem 3.1. For this rate, neither an assumption on the singular system of A nor a signal plus noise assumption such as (17) or (18) is needed.

Theorem 3.2 *Let Assumptions 2.1 and 2.2 be fulfilled, let A fulfill (8), and assume (20). Let $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ and $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ be sequences such that $\lim_{n \rightarrow \infty} \lambda_n = 0$, $\lim_{n \rightarrow \infty} a_n = \infty$,*

$$\lim_{n \rightarrow \infty} \frac{a_n}{\lambda_n^{1+p/2} \sqrt{n}} = 0, \quad \text{and} \quad \lim_{n \rightarrow \infty} \lambda_n a_n = 0. \quad (25)$$

Then,

$$a_n \left(\mathcal{R}_{A,P}(f_{A,\mathbf{D}_n,\lambda_n}) - \mathcal{R}_{A,P}^* \right) \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.} \quad (26)$$

In particular, if $\lambda_n = \gamma n^{-1/(4+p)} \ \forall n \in \mathbb{N}$ for some constant $\gamma \in (0, \infty)$, then every sequence $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ such that

$$\lim_{n \rightarrow \infty} a_n n^{-1/(4+p)} = 0$$

fulfills condition (25).

In the rest of this section, we are concerned with rates of convergence for $f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P}$. While Theorems 2.3, 2.4, 3.1, and 3.2 are valid for all of the commonly used loss functions and do not require any involved assumptions on the distribution P and the operator A , obtaining rates of convergences for $f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P}$ is certainly only possible under much more restrictive and involved assumptions. Therefore, we need some preparations, before we are able to state such rates of convergence in Theorem 3.3 below.

Loss function L .

Nearly all articles on inverse (regression) problems which (at least implicitly) employ a loss function are restricted to the least-squares loss. (A notably exception is Krebs (2011) which uses the ε -insensitive loss common in machine learning.) While we did not have to specify a special loss function so far, we fix a specific loss function now, namely the absolute deviation loss. That is, our loss function is

$$L : \mathcal{Z} \times \mathcal{Y} \times \mathbb{R} \rightarrow [0, \infty), \quad (z, y, t) \mapsto |y - t| \quad (27)$$

in the following. On the one hand, this choice of the loss function is motivated by the fact that the absolute deviation loss typically leads to better robustness properties than the least-squares loss. In particular, it is shown in Hable and Christmann (2011), Christmann and Steinwart (2007), and Christmann et al. (2009) that ordinary regularized kernel methods based on the absolute deviation loss enjoy a qualitative robustness property and have a bounded influence function and a bounded maxbias. On the other hand, the absolute deviation loss together with suitable assumptions on the model guarantee a bound of the form

$$\|g - g^*\|_{L_q(P_Z)} \leq \left(\mathcal{R}_{\text{Id}, P}(g) - \mathcal{R}_{\text{Id}, P}^* \right)^r$$

where $\mathcal{R}_{\text{Id}, P}(g^*) = \int L(z, y, g^*(z)) P(d(z, y)) = \inf_g \int L(z, y, g(z)) P(d(z, y))$ and $q, r \in (0, \infty)$. However, by adapting the model assumptions, such bounds can also be obtained for other loss functions; see (Steinwart and Christmann, 2008, § 3.9).

Inverse regression model.

In the following, we assume the heteroscedastic regression model

$$Y = (Af_{A, P})(Z) + s(Z)\varepsilon, \quad (28)$$

where s is an unknown scale function such that

$$\text{there are constants } \underline{c}_s, \bar{c}_s \in (0, \infty) \text{ with } \underline{c}_s \leq s \leq \bar{c}_s, \quad (29)$$

the random error ε is independent from Z , has

$$\text{median}(\varepsilon) = 0, \quad (30)$$

and

$$\text{the distribution of } \varepsilon \text{ has a Lebesgue-density } h \text{ such that} \quad (31)$$

$$\exists a_h, c_h \in (0, \infty) \text{ s.t. } \forall y \in (-a_h, a_h) : h(y) \geq c_h. \quad (32)$$

If the distribution of the error ε has a Lebesgue-density, then Assumption (32) is very natural and, e.g., fulfilled for all unimodal error distributions. If h is continuous, it is sufficient that $h(0) \neq 0$.

Operator A .

In order to obtain rates of convergences for $f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P}$, we will also need the standard assumption that

$$A : H \rightarrow \mathcal{C}_b(\mathcal{Z}) \text{ is injective.} \quad (33)$$

Let \mathcal{Z}_0 be the support of $P_{\mathcal{Z}}$ and let $P_{\mathcal{Z}_0}$ denote the restriction of $P_{\mathcal{Z}}$ on \mathcal{Z}_0 . Then, the natural embedding $\iota_0 : \mathcal{C}_b(\mathcal{Z}) \rightarrow L_2(P_{\mathcal{Z}_0})$ defines a continuous linear operator $A_0 := \iota_0 \circ A$ and it is easy to see that (33) implies that

$$A_0 : H \rightarrow L_2(P_{\mathcal{Z}_0}) \text{ is injective.} \quad (34)$$

Furthermore, it follows from Assumption (8) and Prop. 6.1 that A_0 is a compact operator. In this case, A_0 has a singular system $(\sigma_j; v_j, u_j)_{j \in \mathbb{N}}$; see, e.g., (Engl et al., 1996, § 2.2). That is, σ_j^2 , $j \in \mathbb{N}$, are the non-zero eigenvalues of the self-adjoint operator $A_0^* A_0$ such that $\sigma_1 \geq \sigma_2 \geq \dots > 0$. The set $\{v_j \mid j \in \mathbb{N}\}$ is a corresponding complete orthonormal system of H ; completeness follows from injectivity of A_0 and (36) below. Finally, the elements

$$u_j := \frac{A_0 v_j}{\|A_0 v_j\|_{L_2(P_{\mathcal{Z}_0})}}, \quad j \in \mathbb{N},$$

form a complete orthonormal system of the closure of $\{A_0 f \mid f \in H\}$ in $L_2(P_{\mathcal{Z}_0})$. We have

$$A_0 v_j = \sigma_j u_j, \quad A_0^* u_j = \sigma_j v_j \quad \forall j \in \mathbb{N} \quad (35)$$

$$A_0 f = \sum_{j=1}^{\infty} \sigma_j \langle f, v_j \rangle_H u_j \quad \forall f \in H \quad (36)$$

$$A_0^* g = \sum_{j=1}^{\infty} \sigma_j \langle g, u_j \rangle_{L_2(P_{\mathcal{Z}_0})} v_j \quad \forall g \in L_2(P_{\mathcal{Z}_0}). \quad (37)$$

Now, we can state the theorem on rates of convergence for $f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P}$.

Theorem 3.3 *Let Assumption 2.1 be fulfilled and assume that the marginal distribution P_Y has a finite first moment, i.e., $\mathbb{E}|Y| = \int |y| P(d(z, y)) < \infty$. Let L be the absolute deviation loss (27) and assume the heteroscedastic model given by (28)–(32). Let A fulfill (8), and (33). Let $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ and $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ be sequences with $\lim_{n \rightarrow \infty} \lambda_n = 0$, $\lim_{n \rightarrow \infty} a_n = \infty$, and*

$$\lim_{n \rightarrow \infty} \frac{\sqrt{a_n}}{\lambda_n \sqrt{n}} = 0. \quad (38)$$

Then, the following assertions are valid:

(a) *If*

$$\sum_{j=1}^{\infty} \frac{|\langle f_{A,P}, v_j \rangle_H|^2}{\sigma_j^2} < \infty \quad (39)$$

and $\lim_{n \rightarrow \infty} a_n \lambda_n^{\frac{1}{2}} = 0$, then

$$a_n \|f_{A, \mathbf{D}_n, \lambda_n} - f_{A,P}\|_H^2 \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.} \quad (40)$$

(b) *Fix any $x \in \mathcal{X}$. If*

$$\sum_{j=1}^{\infty} \frac{(v_j(x))^2}{\sigma_j^2} < \infty \quad (41)$$

and $\lim_{n \rightarrow \infty} a_n \lambda_n = 0$, then

$$a_n (f_{A, \mathbf{D}_n, \lambda_n}(x) - f_{A,P}(x))^2 \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.} \quad (42)$$

For example, in part (a), let $\lambda_n = \gamma n^{-2/5} \forall n \in \mathbb{N}$ for some constant $\gamma \in (0, \infty)$; then all conditions on $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ are fulfilled if $\lim_{n \rightarrow \infty} a_n = \infty$ and

$$\lim_{n \rightarrow \infty} a_n n^{-1/5} = 0.$$

In part (b), let $\lambda_n = \gamma n^{-1/3} \forall n \in \mathbb{N}$ for some constant $\gamma \in (0, \infty)$; then all conditions on $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ are fulfilled if $\lim_{n \rightarrow \infty} a_n = \infty$ and

$$\lim_{n \rightarrow \infty} a_n n^{-1/3} = 0.$$

Assumptions such as (39) in part (a) are often called smoothness assumptions on $f_{A,P}$ and are common in order to obtain rates of convergence. Assumption (41) in part (b) differs as it does not involve $f_{A,P}$ but is a condition on a fixed x . The result of part (b) can be interpreted in the following way: if the inverse regression problem is only moderately ill-posed in some area,

then any $f_{A,P}$ can be estimated with a certain rate of convergence in that area. Of course, the convergence in (40) and (42) could also be reformulated without the exponent 2. However, this formulation makes it possible to easily compare the results with other rates of convergence which most often apply to

$$\mathbb{E} \|f_{A, \mathbf{D}_n, \lambda_n} - f_{A,P}\|_{\star}^2$$

where $\|\cdot\|_{\star}$ denotes the respectively considered norm.

4 Simulation

In order to illustrate the use of regularized kernel methods for inverse regression problems, this section contains simulations in a typical example of an ill-posed problem, namely backward heat conduction. In a statistical setting, this example has also been considered, e.g., in Mair and Ruymgaart (1996) and Bissantz and Holzmann (2013). According to, e.g. (Kress, 1999, Example 15.3), we are faced with the heat equation

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial^2 x}$$

where

$$u : [0, 1] \times [0, T] \rightarrow \mathbb{R}, \quad (x, t) \mapsto u(x, t)$$

denotes the temperature at any spatial point $x \in [0, 1]$ and time $t \in [0, T]$. The boundary conditions are $u(0, t) = u(1, t) = 0$ for every $t \in [0, T]$ and the goal is to recover the initial conditions

$$f(x) := u(x, 0), \quad x \in [0, 1],$$

at time $t = 0$ from (noisy) observations at time $t = T$. Let A denote the operator which maps the initial state f to the final temperature curve $g = u(\cdot, T)$ at time T . Then,

$$(Af)(z) = u(z, T) = \sum_{j=1}^{\infty} \exp(-j^2 \pi^2 T) \int_{[0,1]} f(x) v_j(x) \lambda(dx) v_j(z) \quad (43)$$

with $v_j(z) = \sqrt{2} \sin(j\pi z)$ for every $z \in [0, 1]$. That is $\mathcal{X} = \mathcal{Z} = [0, 1]$ and $(v_j)_{j \in \mathbb{N}}$ is a complete orthonormal system of $L_2([0, 1])$. As a mapping from $L_2([0, 1])$ to $L_2([0, 1])$, the operator A is self-adjoint with eigenfunctions v_j and eigenvalues $\sigma_j = \exp(-j^2 \pi^2 T)$. We use this standard example so that the method can be compared to the spectral cut-off estimator suggested by Mair and Ruymgaart (1996).

The model. We simulate data

$$y_i = (Af_0)(z_i) + s_{\delta}(z_i) \varepsilon_i, \quad i \in \{1, \dots, n\}, \quad (44)$$




Figure 1: The regression function $A(f_0)$ (solid line) and the function f_0 (dotted line) which has to be recovered.

with

$$f_0(x) = -10x(x-1)\sin(4\pi x) \quad \text{and} \quad s_\delta(z) = \delta z(z-1) \quad (45)$$

for $T = 0.01$ and different values of the scale factor δ . The regression function $A(f_0)$ and the function f_0 we want to recover are shown in Figure 1. We use fixed equidistant design points $z_i = (i-1)/n$ on $[0, 1]$ because most theoretical results on spectral cut-off estimators are for such design points. This clearly favors the spectral cut-off estimator also because this estimator is heavily based on the knowledge that the z_i are uniform and the regularized kernel method does not need (and use) such information. The errors ε_i are (independently) sampled from the standard normal distribution. We consider the three scale factors $\delta \in \{0.25, 0.5, 1\}$ which result in errors of similar sizes as in (Bissantz and Holzmann, 2013, § 3.2). The sample sizes are $n \in \{100, 250, 500, 1000\}$.

The estimators. As estimators, we use a regularized kernel method (RKM) and a spectral cut-off estimator (SCE). In case of the RKM, we choose the absolute deviation loss function and the rescaled Wendland kernel

$$k_{1,1}^{(0.3)}(x, x') = \phi_{1,1}(|x - x'|/0.3)$$

with $\phi_{1,1}(r) = (1-r)_+^3(3r+1)$ for every $r \in [0, \infty)$; see (Wendland, 2005, Table 9.1). The scaling 0.3 is approximately equal to the median of the n^2 values $|x_i - x_j|$, $i, j \in \{1, \dots, n\}$. This heuristic is the analogue of a heuristic used in Caputo et al. (2002) and (Karatzoglou et al., 2004, p. 9) in order to choose the scaling factor in case of the Gaussian RBF kernel. The

δ	$n = 100$		$n = 250$		$n = 500$		$n = 1000$	
	RKM	SCE	RKM	SCE	RKM	SCE	RKM	SCE
0.25	0.26	0.32	0.18	0.32	0.14	0.12	0.12	0.09
0.5	0.40	0.32	0.32	0.32	0.24	0.32	0.17	0.16
1	0.46	0.33	0.42	0.32	0.38	0.32	0.30	0.32

Table 2: The medians of the results b_1, \dots, b_{1000} for both estimators in each situation.

regularization parameter is equal to $\lambda = \frac{1}{2}an^{-0.45}$ where a is selected via a 5-fold cross validation among the values

$$10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 5 \cdot 10^{-3}, 10^{-2}, 5 \cdot 10^{-2}, 10^{-1}$$

in each run of the simulation. According to (Mair and Ruymgaart, 1996, § 7.2), the spectral cut-off estimator (SCE) is given by

$$\hat{f}_{n,J}(x) = \sum_{j=1}^J \frac{\hat{b}_{j,n}}{\sigma_j} v_j(x) \quad \text{with} \quad \hat{b}_{j,n} = \frac{1}{n} \sum_{i=1}^n v_j(z_i) y_i$$

where the number $J \in \mathbb{N}$ of basis functions is a regularization parameter which is selected via a 5-fold cross validation among the values

$$2, 3, 4, 5, \dots, 10$$

in each run of the simulation.

Performance results. The simulation consists of 1000 runs. In each run $r \in \{1, \dots, 1000\}$, the quality of the estimate $\hat{f}_n^{(r)}$ is measured by

$$b_r = \int_{[0,1]} |f_0(x) - \hat{f}_n^{(r)}(x)| \lambda(dx) .$$

The medians and the boxplots (the ends of the whiskers represent the 10 and the 90 percent quantiles respectively) of the values b_1, \dots, b_{1000} for both estimators are shown in Table 2 and Figure 2, respectively, in each situation.

In most cases, the performance of both estimators is similar. However, a closer look on the boxplots reveals that, in case of the SCE, the 25, 50, and 75 percent quantiles are very close to each other but the 90 percent quantile is far off for some values of n and δ . The reason for this is that the estimate is very stable for small values of J such as $J = 4$ or $J = 5$ and the 5-fold cross validation most often chooses these values in some situation. However, the SCE turns out to be very sensitive to the choice of J in our simulations. If larger values such as $J \in \{6, 7, 8\}$ are selected, than there is a considerable danger that the SCE breaks down. Therefore, selecting the right value J is

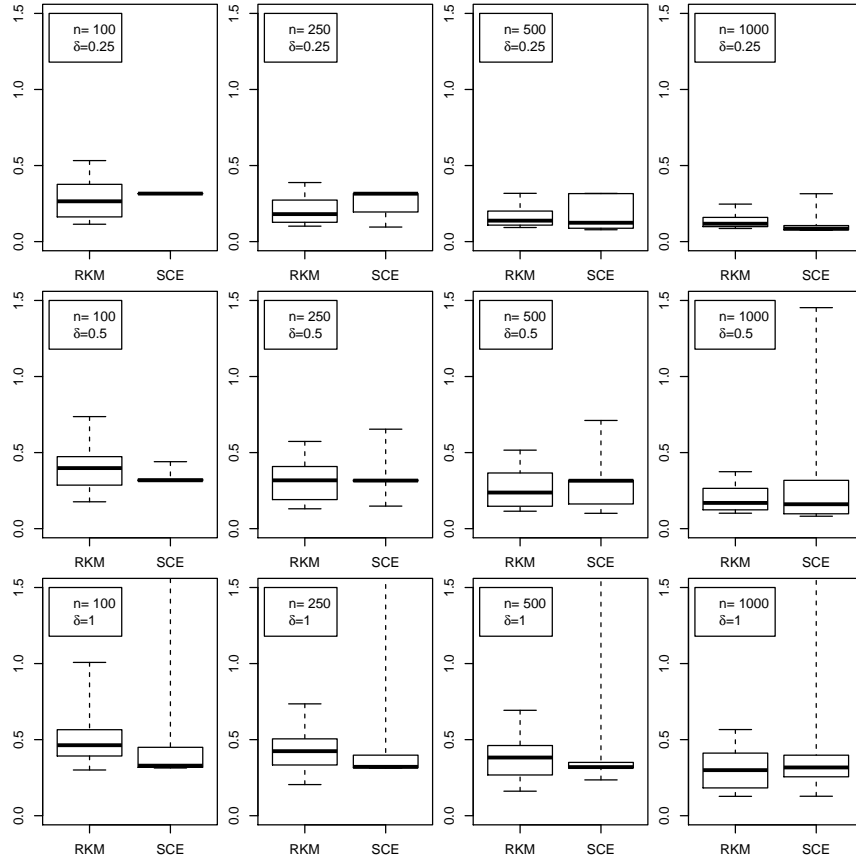


Figure 2: The boxplots of the results b_1, \dots, b_{1000} for both estimators in each situation; the ends of the whiskers represent the 10 and the 90 percent quantiles respectively.

δ	$n = 100$		$n = 250$		$n = 500$		$n = 1000$	
	RKM	SCE	RKM	SCE	RKM	SCE	RKM	SCE
0.25	0.53	0.32	0.39	0.32	0.32	0.32	0.25	0.31
0.5	0.74	0.44	0.57	0.65	0.52	0.71	0.37	1.45
1	1.01	5.33	0.74	4.89	0.69	2.68	0.57	5.59

Table 3: The 90 percent quantiles of the results b_1, \dots, b_{1000} for both estimators in each situation.

crucial and automatic data-driven methods such as a k -fold cross validation might not be sufficient. Table 3 shows the 90 percent quantiles of the values b_1, \dots, b_{1000} for both estimators in each situation. From these quantiles, it can be seen that the SCE frequently breaks down for larger values of n and δ . Therefore, Table 2 does not show the mean but the median of the values b_1, \dots, b_{1000} because the mean is corrupted by large outliers in case of the SCE. As this does not happen in case of the RKM, using the median favors the SCE.

Details on computations. The computation of the spectral cut-off estimator is simple and extremely fast in this case because the spectral decomposition is already known here. As mentioned below Theorem 2.3, the computation of regularized kernel methods for inverse problems can be done by using standard software for computing ordinary regularized kernel methods. The only additional thing one has to do is to calculate the pseudo kernel matrix M as defined in (15). Here, we have

$$M_{i,j} = \sum_{q=1}^{\infty} \sum_{s=1}^{\infty} \exp(-(q^2 + s^2)\pi^2 T) \mathcal{I}_{q,s} v_q(z_j) v_s(z_i)$$

for

$$\mathcal{I}_{q,s} = \int_{[0,1]^2} k_{1,1}^{(0,3)}(x_1, x_2) v_q(x_1) v_s(x_2) \lambda^2(d(x_1, x_2)).$$

As the exponential coefficients decrease extremely fast, the infinite double series can be approximated very well by only calculating a few terms, e.g., up to $q, s = 30$ is more than enough. The double integrals $\mathcal{I}_{q,s}$ are approximated by a Monte-Carlo simulation in our simulated example. Then, the coefficients $\alpha_1, \dots, \alpha_n$ are calculated by the R-package “kernlab” (a standard software for calculating ordinary regularized kernel methods), in the following way:

```
model ← kqr(as.kernelMatrix(M), y, tau=0.5, C=cost)
alpha ← alpha(model)
```

where y denotes the vector of the observed values y_i , $\text{tau}=0.5$ corresponds to using the absolute deviation loss function, and cost is the cost regularization

parameter which is equal to $1/(2n\lambda)$. It is worth mentioning that it is sufficient to calculate M once and to reuse this matrix in every step of the k -fold cross validation for tuning `cost`. Finally, values of the estimate \hat{f}_n can be calculated according to (14); here, we have

$$\hat{f}_n(x) = \sum_{q=1}^{\infty} \exp(-q^2 \pi^2 T) \mathcal{I}_q(x) \sum_{i=1}^n \alpha_i v_q(z_i)$$

for

$$\mathcal{I}_q(x) = \int_{[0,1]} k_{1,1}^{(0,3)}(x_1, x) v_q(x_1) \lambda(dx_1).$$

Again, the extremely fast decay of the exponential coefficients guarantees that the series can be approximated very well by only calculating a few terms.

5 Conclusions

Regularized kernel methods constitute a broad and flexible class of methods which originate from machine learning and are common in nonparametric classification and regression problems today. In this article, we investigated the use of regularized kernel methods for inverse problems in a unifying way. In addition to consistency results under very weak assumptions, we also obtained a rate of convergence under a typical smoothness assumption on the target function. Though such a rate of convergence is of a purely theoretical manner and is not interesting on its own for real data analysis, it can play an important role in developing methods for statistical inference (such as asymptotic confidence sets) based on undersmoothing. However, statistical inference is still at an early stage even in case of regularized kernel methods for ordinary regression problems as well as in case of inverse regression problems with any other estimation method. First steps on statistical inference for regularized kernel methods are done in De Brabanter et al. (2011), Hable (2012b), and Hable (2012a). In case of inverse regression problems, Bissantz and Holzmann (2008) and Bissantz and Birke (2009) are concerned with asymptotic confidence sets for spectral cut-off estimators. Accordingly, enabling statistical inference for inverse regression problems with regularized kernel methods is a matter of future and challenging research. Using regularized kernel methods in real data analysis of inverse regression problems seems to be promising as they have nice computational properties and it is possible to resort to already existing well developed software implementations of ordinary regularized kernel methods.

6 Appendix

6.1 Additional Results

This subsection contains some additional results on regularized kernel methods for inverse problems. On the one hand, these results are needed in the proofs of the main results; on the other hand, they are also interesting of its own, in particular, as some of them are the counterparts of some of the main tools for ordinary regularized kernel methods (i.e. $A = \text{id}$).

The first proposition shows that, in our setting, compactness of A (in a rather weak sense) comes for free by assuming (8).

Proposition 6.1 *Let Assumption 2.1 be fulfilled, and assume that A fulfills (8). Then, $A : H \rightarrow \mathcal{C}_b(\mathcal{Z})$ is a compact operator; that is,*

$$(f_n)_{n \in \mathbb{N}_0} \subset H, \quad f_n \xrightarrow{w} f_0 \quad \Rightarrow \quad \|A(f_n) - A(f_0)\|_\infty \longrightarrow 0 \quad (46)$$

where \xrightarrow{w} denotes weak convergence in the Hilbert space.

The next theorem is a general representer theorem. In case of ordinary regularized kernel methods, general representer theorems are *the* main tool for deriving theoretical results on consistency, rates of convergence, asymptotic normality, and robustness. The proof of the following theorem for the case of inverse problems is similar to the proof of the general representer theorem in the ordinary case (Theorem 5.8 and Theorem 5.9 Steinwart and Christmann, 2008, see, e.g.,) even though the assumptions and the result considerably differ.

Theorem 6.2 (General Representer Theorem)

Let Assumptions 2.1 and 2.2 be fulfilled, and fix any $\lambda > 0$. Then, there is an $h_{P,\lambda} \in \mathcal{L}_2(P)$ with the following properties:

(a) *For every $(z, y) \in \mathcal{Z} \times \mathcal{Y}$,*

$$|h_{P,\lambda}(z, y)| \leq b'_0(y) + b'_1 \cdot \left(\|A\| \sqrt{\frac{1}{\lambda} \int b \, dP} + 1 \right)^p, \quad (47)$$

$$h_{P,\lambda}(z, y) \in \partial L(z, y, (Af_{A,P,\lambda})(z)) \quad (48)$$

where $\partial L(z, y, \cdot)$ denotes the subdifferential of the convex function $t \mapsto L(z, y, t)$.

(b) *If $A_P^* : L_2(P) \rightarrow H$ denotes the adjoint of the continuous linear map $A_P : H \rightarrow L_2(P)$ given by $(A_P f)(z, y) = (Af)(z) \quad \forall f \in H, z \in \mathcal{Z}, y \in \mathcal{Y}$, then*

$$f_{A,P,\lambda} = -\frac{1}{2\lambda} A_P^*(h_{P,\lambda}) \quad (49)$$

and

$$f_{A,P,\lambda}(x) = -\frac{1}{2\lambda} \int A(\Phi(x))(z) h_{P,\lambda}(z, y) P(d(z, y)) \quad \forall x \in \mathcal{X}. \quad (50)$$

(c) If P_1 is a probability measure on $\mathcal{Z} \times \mathcal{Y}$ such that $\int b dP_1 < \infty$ and $\int b_0^2 dP_1 < \infty$, then

$$\|f_{A,P_1,\lambda} - f_{A,P,\lambda}\|_H \leq \frac{1}{\lambda} \sup_{f \in \mathcal{F}} \left| \int h_{P,\lambda} A f dP_1 - \int h_{P,\lambda} A f dP \right| \quad (51)$$

where $\mathcal{F} = \{f \in H \mid \|f\|_H \leq 1\}$.

The following theorem yields a rate of convergence for the stochastic part, that is, the difference between the empirical estimate $f_{A,\mathbf{D}_n,\lambda_n}$ and its theoretical counterpart f_{A,P,λ_n} . In case of ordinary regularized kernel methods (i.e. $A = \text{id}$), a corresponding result can simply be proven by the representer theorem and Hoeffding's inequality for Hilbert spaces. However, in our case of inverse problems, the situation is much more complicated because working with property (51) of the representer theorem for inverse problems is more troublesome than working with the corresponding property in case of $A = \text{id}$. Accordingly, we cannot apply Hoeffding's inequality offhand but use Donsker theory for empirical processes instead.

Theorem 6.3 *Let Assumptions 2.1 and 2.2 be fulfilled. Let $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ and $(a_n)_{n \in \mathbb{N}} \subset (0, \infty)$ be sequences such that*

$$\lim_{n \rightarrow \infty} \lambda_n = 0, \quad \lim_{n \rightarrow \infty} a_n = \infty, \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{a_n}{\lambda_n^{1+p/2} \sqrt{n}} = 0. \quad (52)$$

Let (8) be fulfilled for A . Then,

$$a_n \|f_{A,\mathbf{D}_n,\lambda_n} - f_{A,P,\lambda_n}\|_H \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.}$$

In the following, we are concerned with the deterministic part. Prop. 6.4 states that the risk of f_{A,P,λ_n} converges to the infimal risk; Theorem 6.5 yields that f_{A,P,λ_n} even converges in H -norm to a minimizer of the risk – provided that a minimizer exists in H .

Proposition 6.4 *Let Assumption 2.1 and (19) be fulfilled, and let $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ be a sequence such that $\lim_{n \rightarrow \infty} \lambda_n = 0$. Then,*

$$\lim_{n \rightarrow \infty} \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) = \inf_{f \in H} \mathcal{R}_{A,P}(f).$$

Theorem 6.5 *Let Assumptions 2.1 and 2.2 be fulfilled, and let A fulfill (8). Assume that*

$$\exists f^* \in H \text{ s.t. } \mathcal{R}_{A,P}(f^*) = \inf_{f \in H} \mathcal{R}_{A,P}(f). \quad (53)$$

Then, there is a unique $f_{A,P} \in H$ with the following two properties:

$$\mathcal{R}_{A,P}(f_{A,P}) = \inf_{f \in H} \mathcal{R}_{A,P}(f), \quad (54)$$

$$f^* \in H, \mathcal{R}_{A,P}(f^*) = \inf_{f \in H} \mathcal{R}_{A,P}(f) \Rightarrow \|f^*\|_H > \|f_{A,P}\|_H \text{ or } f^* = f_{A,P}. \quad (55)$$

Furthermore, for every sequence $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ such that $\lim_{n \rightarrow \infty} \lambda_n = 0$, it follows that

$$\lim_{n \rightarrow \infty} \|f_{A,P,\lambda_n} - f_{A,P}\|_H = 0. \quad (56)$$

The following Lemma 6.6 provides us with a bound of the form (28). Such a bound could also easily be adopted from the general results in (Steinwart and Christmann, 2008, § 3.9). However, in our special situation, it is possible to obtain a tighter bound which enables the proof of better rates of convergence in Theorem 3.3.

Lemma 6.6 *Let P be a probability measure on $\mathcal{Z} \times \mathcal{Y}$ such that the marginal distribution $P_{\mathcal{Y}}$ has a finite first moment, i.e., $\int |y| P(dy) < \infty$. Let L be the absolute deviation loss (27) and assume the heteroscedastic model given by (28)–(32). Define $\alpha := a_h \underline{c}_s > 0$ and $t_z^* := (Af_{A,P})(z)$ for every $z \in \mathcal{Z}$. Then, there is a $B \in \mathfrak{B}_{\mathcal{Z}}$ such that $P_{\mathcal{Z}}(B) = 1$ and, for every $z \in B$ and $t \in (t_z^* - \alpha, t_z^* + \alpha)$,*

$$\frac{c_h}{2\bar{c}_s} \cdot (t - t_z^*)^2 \leq \int L(z, y, t) P(dy|z) - \int L(z, y, t_z^*) P(dy|z).$$

6.2 Proofs

Proof of Theorem 2.3: Fix any $D_n = ((z_1, y_1), \dots, (z_n, y_n)) \in (\mathcal{Z} \times \mathcal{Y})^n$ and $\lambda \in (0, \infty)$. Existence and uniqueness of $f_{A,D_n,\lambda}$ defined by (13) follow from Theorem 2.4 by choosing the empirical measure for P . Define

$$\tilde{A} : H \rightarrow \mathbb{R}^n, \quad f \mapsto \left((Af)(z_1), \dots, (Af)(z_n) \right)^\top.$$

The assumptions on A imply that \tilde{A} is again linear and continuous. Let \tilde{A}^* denote the adjoint operator of \tilde{A} . Then, for every $x \in \mathcal{X}$,

$$(\tilde{A}^*(e_i))(x) = \langle \tilde{A}^*(e_i), \Phi(x) \rangle_H = \langle e_i, \tilde{A}(\Phi(x)) \rangle_{\mathbb{R}^n} = \left(A(\Phi(x)) \right)(z_i) \quad (57)$$

and, therefore,

$$\begin{aligned} M_{i,j} &= \left(A \left((A\Phi(\cdot))(z_i) \right) \right)(z_j) \stackrel{(57)}{=} \left(A(\tilde{A}^*(e_i)) \right)(z_j) = \left\langle e_j, \tilde{A}(\tilde{A}^*(e_i)) \right\rangle_{\mathbb{R}^n} \\ &= \langle \tilde{A}^*(e_j), \tilde{A}^*(e_i) \rangle_H \end{aligned} \quad (58)$$

This implies that the matrix M is symmetric and, in addition, that it is positive semi-definite because, for every $\alpha = (\alpha_1, \dots, \alpha_n)^\top \in \mathbb{R}^n$,

$$\alpha^\top M \alpha \stackrel{(58)}{=} \sum_{i,j} \alpha_i \alpha_j \langle \tilde{A}^*(e_i), \tilde{A}^*(e_j) \rangle_H = \left\| \sum_i \alpha_i \tilde{A}^*(e_i) \right\|_H^2 \geq 0. \quad (59)$$

Fix any $\alpha = (\alpha_1, \dots, \alpha_n)^\top \in \mathbb{R}^n$ and define

$$f_0(x) = \sum_{i=1}^n \alpha_i \cdot (A\Phi(x))(z_i) \quad \forall x \in \mathcal{X}. \quad (60)$$

Note that (57) implies that $f_0 \in H$ and

$$\|f_0\|_H^2 \stackrel{(57)}{=} \left\| \sum_i \alpha_i \tilde{A}^*(e_i) \right\|_H^2 \stackrel{(59)}{=} \alpha^\top M \alpha \quad (61)$$

Furthermore,

$$\begin{aligned} (Af_0)(z_j) &= \langle e_j, \tilde{A}(f_0) \rangle_{\mathbb{R}^n} \stackrel{(60,57)}{=} \sum_{i=1}^n \alpha_i \langle e_j, \tilde{A}(\tilde{A}^*(e_i)) \rangle_{\mathbb{R}^n} = \\ &\stackrel{(58)}{=} \sum_{i=1}^n \alpha_i M_{i,j} = \alpha^\top M e_j. \end{aligned} \quad (62)$$

Hence, (16) follows from the definition of the regularized risk $\mathcal{R}_{A,D,\lambda}$, (61), and (62).

It only remains to prove (14), which can be done similarly to the proof of (Krebs et al., 2009, Theorem 3.2) and (Krebs, 2011, Lemma 3.1). The main idea of the proof is to show that $f_{A,D_n,\lambda}$ is an element of the image $\text{im}(\tilde{A}^*)$. First, note that the image $\text{im}(\tilde{A}^*)$ is a finite-dimensional linear subspace of H , hence, it is closed; see, e.g. (Denkowski et al., 2003, Cor. 3.2.17). Then, for every $f_0 \in H$, there is an $f_1 \in \text{im}(\tilde{A}^*)$ and an $f_2 \in (\text{im}(\tilde{A}^*))^\perp$ such that $f_0 = f_1 + f_2$; see, e.g., (Denkowski et al., 2003, Cor. 3.7.16). Then,

$$\begin{aligned} (Af_0)(z_i) &= (Af_1)(z_i) + (Af_2)(z_i) = (Af_1)(z_i) + \langle e_i, \tilde{A}(f_2) \rangle_H = \\ &= (Af_1)(z_i) + \langle \tilde{A}^*(e_i), f_2 \rangle_H = (Af_1)(z_i) \end{aligned}$$

and

$$\|f_0\|_H^2 = \|f_1\|_H^2 + \|f_2\|_H^2.$$

This shows that, if f_0 is not in the image of \tilde{A}^* , then there is another $f_1 \in H$ such that $\mathcal{R}_{A,D,\lambda}(f_1) < \mathcal{R}_{A,D,\lambda}(f_0)$. Hence, the minimizer $f_{A,D_n,\lambda}$ is in the image of \tilde{A}^* , that is, there are $\alpha_1, \dots, \alpha_n \in \mathbb{R}$ such that

$$f_{A,D_n,\lambda} = \tilde{A}^* \left(\sum_{i=1}^n \alpha_i e_i \right) = \sum_{i=1}^n \alpha_i \tilde{A}^*(e_i) \stackrel{(57)}{=} \sum_{i=1}^n \alpha_i \cdot (A\Phi(\cdot))(z_i).$$

□

Proof of Theorem 2.4: Consider $\mathcal{R}_{A,P,\lambda}$ as a map from H to $\mathbb{R} \cup \{\pm\infty\}$. the map $\mathcal{R}_{A,P,\lambda}$ is convex and fulfills $\mathcal{R}_{A,P,\lambda}(f) \geq 0 > -\infty$ for every $f \in H$ and $\lim_{\|f\|_H \rightarrow \infty} \mathcal{R}_{A,P,\lambda}(f) = \infty$ because

$$\liminf_{\|f\|_H \rightarrow \infty} \mathcal{R}_{A,P,\lambda}(f) \geq \lambda \liminf_{\|f\|_H \rightarrow \infty} \|f\|_H^2 = \infty.$$

According to Assumption (7),

$$f_n \xrightarrow[n \rightarrow \infty]{} f \quad \text{in } H \quad \Rightarrow \quad (Af_n)(z) \xrightarrow[n \rightarrow \infty]{} (Af)(z) \quad \forall z \in \mathcal{Z}.$$

so that Fatou's Lemma – e.g. (Denkowski et al., 2003, Theorem 2.2.17) – implies that $\mathcal{R}_{A,P,\lambda}$ is lower semicontinuous. Then, it follows from (Denkowski et al., 2003, Prop. 5.2.12) that there is an $f_{A,P,\lambda} \in H$ which minimizes $\mathcal{R}_{A,P,\lambda}$ in H . Assumption (19) implies $\mathcal{R}_{A,P}(f_{A,P,\lambda}) < \infty$ so that uniqueness of $f_{A,P,\lambda} \in H$ follows from strict convexity of the squared norm and convexity of $\mathcal{R}_{A,P}$. \square

Proof of Prop. 6.1: See, e.g., (Denkowski et al., 2003, Prop. 3.7.47) for the fact that compactness of $A : H \rightarrow \mathcal{C}_b(\mathcal{Z})$ is equivalent to (46). In order to show (46), fix any sequence $(f_n)_{n \in \mathbb{N}} \subset H$ which converges weakly in H to some $f_0 \in H$ for $n \rightarrow \infty$. As a weakly convergent sequence in a Hilbert space is bounded (see, e.g., Denkowski et al., 2003, Cor. 3.4.10), there is a $c \in (0, \infty)$ such that, for every $n \in \mathbb{N}$, we have $\|f_n\|_H \leq c$. Hence, for every sequence $x_\ell \rightarrow x_0$ in \mathcal{X} ,

$$\begin{aligned} \limsup_{\ell \rightarrow \infty} \sup_{n \in \mathbb{N}} |f_n(x_\ell) - f_n(x_0)| &\stackrel{(10)}{=} \limsup_{\ell \rightarrow \infty} \sup_{n \in \mathbb{N}} |\langle f_n, \Phi(x_\ell) - \Phi(x_0) \rangle_H| \leq \\ &\leq \limsup_{\ell \rightarrow \infty} \sup_{n \in \mathbb{N}} \|f_n\|_H \cdot \|\Phi(x_\ell) - \Phi(x_0)\|_H \leq c \cdot \lim_{\ell \rightarrow \infty} \|\Phi(x_\ell) - \Phi(x_0)\|_H = 0 \end{aligned}$$

since continuity of k implies continuity of Φ ; see, e.g., (Steinwart and Christmann, 2008, Lemma 4.29). That is, we have shown that the sequence $(f_n)_{n \in \mathbb{N}}$ is equicontinuous. In addition, it follows from weak convergence in H and the reproducing property (10) that f_n converges to f_0 pointwise. Since \mathcal{X} is compact, pointwise convergence together with equicontinuity implies uniform convergence of f_n to f_0 ; see, e.g., (Denkowski et al., 2003, Prop. 1.6.14 and Theorem 1.6.12). Hence, the statement follows from assumption (8). \square

Proof of Theorem 6.2: We start with the proof of (a) and (b). Define

$$\mathcal{Q}_P : L_2(P) \rightarrow \mathbb{R}, \quad g \mapsto \int L(z, y, g(z, y)) P(d(z, y)).$$

That is, the risk $\mathcal{R}_{A,P} : H \rightarrow \mathbb{R}$ is given by $\mathcal{R}_{A,P}(f) = (\mathcal{Q}_P \circ A_P)(f)$, $f \in H$. Note that Assumption 2.2 implies that \mathcal{Q}_P is defined well. Then,

it follows from (Rockafellar, 1976, Prop. 2C and Cor. 3E) that the subdifferential of the convex map \mathcal{Q}_P is given by

$$\partial \mathcal{Q}_P(g) = \left\{ h \in L_2(P) \mid h(z, y) \in \partial L(z, y, g(z, y)) \quad P(d(z, y)) - \text{a.s.} \right\};$$

see, e.g., (Steinwart and Christmann, 2008, Prop. A.6.13). It is easy to see from Assumption 2.2 that \mathcal{Q}_P is continuous on $L_2(P)$ and, therefore,

$$\partial \mathcal{R}_{A,P}(f) = \partial(\mathcal{Q}_P \circ A_P)(f) = A_P^*(\partial \mathcal{Q}_P(A_P f));$$

see, e.g., (Denkowski et al., 2003, Theorem 5.3.33). That is,

$$\partial \mathcal{R}_{A,P}(f) = \left\{ A_P^*(h) \mid \begin{array}{l} h \in L_2(P), \\ h(z, y) \in \partial L(z, y, (A_P f)(z, y)) \quad P(d(z, y)) - \text{a.s.} \end{array} \right\}.$$

The convex map $H \rightarrow \mathbb{R}$, $f \mapsto \lambda \|f\|_H^2$ is Fréchet differentiable with derivative $2\lambda f$ – see, e.g., (Denkowski et al., 2003, Example 5.1.6(c)) – and, therefore, its subdifferential at $f \in H$ is given by $\{2\lambda f\}$ – see, e.g., (Denkowski et al., 2003, Prop. 5.3.30). Since $\mathcal{R}_{A,P,\lambda}(f) = \mathcal{R}_{A,P}(f) + \lambda \|f\|_H^2$ for every $f \in H$, it follows that

$$\partial \mathcal{R}_{A,P,\lambda}(f) = 2\lambda f + \partial \mathcal{R}_{A,P}(f), \quad f \in H;$$

see, e.g., (Denkowski et al., 2003, Theorem 5.3.32). Since $\mathcal{R}_{A,P,\lambda}$ attains its minimum in H at $f_{A,P,\lambda}$, it follows from the definition of the subdifferential that $0 \in \partial \mathcal{R}_{A,P,\lambda}(f_{A,P,\lambda})$. That is, there is an $h \in \mathcal{L}_2(P)$ such that $h(z, y) \in \partial L(z, y, (A_P f)(z, y))$ for P -a.e. $(z, y) \in \mathcal{Z} \times \mathcal{Y}$ and $f_{A,P,\lambda} = -\frac{1}{2\lambda} A_P^*(h)$. In the following, it is shown that we can even choose $h_{P,\lambda} \in \mathcal{L}_2(P)$ such that $h_{P,\lambda}(z, y) \in \partial L(z, y, (A_P f)(z, y))$ for every $(z, y) \in \mathcal{Z} \times \mathcal{Y}$. For every $(z, y) \in \mathcal{Z} \times \mathcal{Y}$, let $L'_+(z, y, \cdot)$ denote the right derivative function of $L(z, y, \cdot)$. Recall from (Rockafellar, 1970, Theorem 24.1 and p. 229), that this is a function $L'_+(z, y, \cdot) : \mathbb{R} \rightarrow \mathbb{R}$ and $L'_+(z, y, t) \in \partial L(z, y, t)$ for every $(z, y) \in \mathcal{Z} \times \mathcal{Y}$ and $t \in \mathbb{R}$. Since the function $L'_+ : (z, y, t) \mapsto L'_+(z, y, t)$ is the pointwise limit of a sequence of measurable functions, L'_+ is measurable. Hence, there is a P -null-set $N \in \mathfrak{B}_{\mathcal{Z} \times \mathcal{Y}}$ such that $h_{P,\lambda} \in \mathcal{L}_2(P)$ defined by

$$h_{P,\lambda}(z, y) = h(z, y) I_{N^c}(z, y) + L'_+(z, y, (A_P f)(z, y)) I_N(z, y)$$

fulfills (48) and (49). Next, (50) follows from

$$\begin{aligned} f_{A,P,\lambda}(x) &\stackrel{(10)}{=} \langle f_{A,P,\lambda}, \Phi(x) \rangle_H \stackrel{(49)}{=} -\frac{1}{2\lambda} \langle A_P^* h_{P,\lambda}, \Phi(x) \rangle_H = \\ &= -\frac{1}{2\lambda} \langle h_{P,\lambda}, A_P(\Phi(x)) \rangle_{L_2(P)} = -\frac{1}{2\lambda} \int A(\Phi(x)) h_{P,\lambda} dP \end{aligned}$$

for every $x \in \mathcal{X}$. Finally, in order to prove (47), note that

$$\lambda \|f_{A,P,\lambda}\|_H^2 \leq \mathcal{R}_{A,P,\lambda}(f_{A,P,\lambda}) \leq \mathcal{R}_{A,P,\lambda}(0) \stackrel{(11)}{\leq} \int b dP$$

and, therefore, (7) implies

$$\|Af_{A,P,\lambda}\|_\infty \leq \|A\| \cdot \|f_{A,P,\lambda}\|_H < \|A\| \sqrt{\frac{1}{\lambda} \int b dP} + 1 =: a < \infty. \quad (63)$$

Now, fix any $(z, y) \in \mathcal{Z} \times \mathcal{Y}$, $t \in (-a, a)$, and $\gamma_t \in \partial L(z, y, t)$. Then, the definition of the subdifferential implies

$$\gamma_t \cdot (a - t) \leq L(z, y, a) - L(z, y, t) \stackrel{(12)}{\leq} (b'_0(y) + b'_1 a^p) \cdot |a - t|$$

and

$$\gamma_t \cdot (-a - t) \leq L(z, y, -a) - L(z, y, t) \stackrel{(12)}{\leq} (b'_0(y) + b'_1 a^p) \cdot |-a - t|.$$

Dividing these inequalities by $|a - t| = a - t$ and $-|-a - t| = -a - t$, respectively, leads to

$$|\gamma_t| \leq b'_0(y) + b'_1 a^p. \quad (64)$$

Due to (63), we may choose $t = (Af_{A,P,\lambda})(z) \in (-a, a)$ so that (48) and (64) yield

$$|h_{P,\lambda}(z, y)| \leq b'_0(y) + b'_1 a^p.$$

Then, (47) follows from the definition of a in (63); that is, we have proven parts (a) and (b) of the theorem.

For the proof of part (c), let P_1 be any probability measure on $\mathcal{Z} \times \mathcal{Y}$ such that $\int b dP_1 < \infty$ and $\int b_0'^2 dP_1 < \infty$. In order to shorten the notation, define $h_0 := h_{P,\lambda}$, $f_0 := f_{A,P,\lambda}$, and $f_1 := f_{A,P_1,\lambda}$. Then, (48) implies

$$h_0(z, y)(Af_1(z) - Af_0(z)) \leq L(z, y, Af_1(z)) - L(z, y, Af_0(z)) \quad (65)$$

for every $(z, y) \in \mathcal{Z} \times \mathcal{Y}$. The map $A_{P_1} : H \rightarrow L_2(P_1)$ defined by $(A_{P_1}f)(z, y) = (Af)(z)$ is a continuous linear operator; let $A_{P_1}^*$ denote its adjoint operator. Since $h_0 \in L_2(P_1)$ according to (47) and the assumptions on P_1 , it follows that

$$\begin{aligned} \int h_0(z, y)(Af_1(z) - Af_0(z)) P_1(d(z, y)) &= \langle h_0, A_{P_1}(f_1 - f_0) \rangle_{L_2(P_1)} = \\ &= \langle f_1 - f_0, A_{P_1}^* h_0 \rangle_H. \end{aligned}$$

Hence, integrating both sides of (65) with respect to P_1 implies

$$\langle f_1 - f_0, A_{P_1}^* h_0 \rangle_H \leq \mathcal{R}_{A,P_1}(f_1) - \mathcal{R}_{A,P_1}(f_0). \quad (66)$$

An elementary calculation shows

$$2\lambda \langle f_1 - f_0, f_0 \rangle_H + \lambda \|f_0 - f_1\|_H^2 = \lambda \|f_1\|_H^2 - \lambda \|f_0\|_H^2. \quad (67)$$

Then, calculating (66) + (67), the definition of the regularized risk $\mathcal{R}_{A, P_1, \lambda}$, and the definition of $f_1 = f_{A, P_1, \lambda}$ imply

$$\langle f_1 - f_0, A_{P_1}^* h_0 + 2\lambda f_0 \rangle_H + \lambda \|f_0 - f_1\|_H^2 \leq \mathcal{R}_{A, P_1, \lambda}(f_1) - \mathcal{R}_{A, P_1, \lambda}(f_0) \leq 0.$$

Hence, it follows from (49) that

$$\lambda \|f_0 - f_1\|_H^2 \leq \langle f_0 - f_1, A_{P_1}^* h_0 - A_P^* h_0 \rangle_H \leq \|f_0 - f_1\|_H \|A_{P_1}^* h_0 - A_P^* h_0\|_H$$

and, therefore,

$$\begin{aligned} \lambda \|f_1 - f_0\|_H &\leq \|A_{P_1}^* h_0 - A_P^* h_0\|_H = \sup_{f \in \mathcal{F}} |\langle f, A_{P_1}^* h_0 \rangle_H - \langle f, A_P^* h_0 \rangle_H| \\ &= \sup_{f \in \mathcal{F}} \left| \langle A_{P_1} f, h_0 \rangle_{L_2(P_1)} - \langle A_P f, h_0 \rangle_{L_2(P)} \right| = \\ &= \sup_{f \in \mathcal{F}} \left| \int h_0(z, y) (Af)(z) P_1(d(z, y)) - \int h_0(z, y) (Af)(z) P(d(z, y)) \right|. \end{aligned}$$

□

Proof of Theorem 6.3: For every $n \in \mathbb{N}$, take $h_n := h_{P, \lambda_n}$ from Theorem 6.2 and define

$$g_{n, f} := \frac{a_n}{\lambda_n \sqrt{n}} h_n A f \quad \forall f \in \mathcal{F}$$

where $\mathcal{F} = \{f \in H \mid \|f\|_H \leq 1\}$. Then,

$$\mathcal{G}_n := \{g_{n, f} \mid f \in \mathcal{F}\}$$

is a changing class in the sense of (van der Vaart, 1998, § 19.5) and we can use a Donsker theorem for such classes in order to prove

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n g_{n, f}(Z_i, Y_i) - \mathbb{E}_P g_{n, f} \right)_{f \in \mathcal{F}} \rightsquigarrow 0 \quad \text{in } \ell_\infty(\mathcal{F}) \quad (68)$$

below. Then, it follows from (68), the definition of $g_{n, f}$, and the continuous mapping theorem (e.g., van der Vaart and Wellner, 1996, Theorem 1.3.6) that

$$\frac{a_n}{\lambda_n} \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n h_{P, \lambda_n}(Z_i, Y_i) A f(Z_i) - \int h_{P, \lambda_n} A f dP \right| \rightsquigarrow 0. \quad (69)$$

Since weak convergence to a constant implies convergence in (outer) probability (see, e.g., van der Vaart and Wellner, 1996, Lemma 1.10.2), the

statement of Theorem 6.3 follows from (69) and (51). Measurability of the random variable $\|f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P, \lambda_n}\|_H$ can be proven by simply following the lines of the proof of (Hable, 2013, Lemma 9a). (The only noteworthy difference is that proving the analogue of (Hable, 2013, (54)) involves an application of Prop. 6.1.)

Hence, it only remains to prove (68) and this is done by use of (van der Vaart, 1998, § 19.5) in the following. To this end, note that $\mathcal{G}_n = \{g_{n,f} \mid f \in \mathcal{F}\}$ is a class of measurable functions indexed by \mathcal{F} . For the metric ρ defined by $\rho(f_1, f_2) = \|f_1 - f_2\|_\infty$, $f_1, f_2 \in \mathcal{F}$, the index set \mathcal{F} is totally bounded because \mathcal{F} is relatively compact in $\mathcal{C}_b(\mathcal{X})$ according to (Steinwart and Christmann, 2008, Cor. 4.31). It follows from $\lim_{n \rightarrow \infty} \lambda_n = 0$ and (47) that there is a $c \in (0, \infty)$ such that, for every $n \in \mathbb{N}$,

$$|h_n(z, y)| \leq b'_0(y) + c\lambda_n^{-p/2} \quad \forall (z, y) \in \mathcal{Z} \times \mathcal{Y}. \quad (70)$$

Since $\|Af_1 - Af_2\|_\infty \leq \|A\| \cdot \|f_1 - f_2\|_H \leq 2\|A\|$ for every $f_1, f_2 \in \mathcal{F}$, it follows from the definitions that

$$\int (g_{n,f_1} - g_{n,f_2})^2 dP \leq 4\|A\|^2 \cdot \left(\frac{a_n}{\lambda_n \sqrt{n}}\right)^2 \int h_n^2 dP \quad \forall f_1, f_2 \in \mathcal{F}. \quad (71)$$

Then, an easy calculation using (52), (70), and (71) shows

$$\lim_{n \rightarrow \infty} \sup_{f_1, f_2 \in \mathcal{F}} \int (g_{n,f_1} - g_{n,f_2})^2 dP = 0. \quad (72)$$

Due to (52), there is a $G \in \mathcal{L}_2(P_{\mathcal{Y}})$ such that, for every $n \in \mathbb{N}$ and $(z, y) \in \mathcal{Z} \times \mathcal{Y}$,

$$\frac{a_n \|A\|}{\lambda_n \sqrt{n}} |h_n(z, y)| \stackrel{(70)}{\leq} \frac{a_n \|A\|}{\lambda_n \sqrt{n}} b'_0(y) + \frac{a_n \|A\|}{\lambda_n^{1+p/2} \sqrt{n}} c \leq G(y). \quad (73)$$

Since $\|Af\|_\infty \leq \|A\| \cdot \|f\|_H \leq \|A\|$ for every $f \in \mathcal{F}$, it follows from (73) that

$$|g_{n,f}(z, y)| \leq G(y) \quad \forall (z, y) \in \mathcal{Z} \times \mathcal{Y}, \quad n \in \mathbb{N}, \quad f \in \mathcal{F}. \quad (74)$$

Hence, G is an envelope function of \mathcal{G}_n for every $n \in \mathbb{N}$ which fulfills the Lindeberg conditions $\int G^2 dP < \infty$ and $\lim_{n \rightarrow \infty} \int G^2 I_{G > \varepsilon \sqrt{n}} dP = 0$ for every $\varepsilon > 0$.

Let $\|A\|_\infty$ denote the operator norm of A as an operator from $\mathcal{C}_b(\mathcal{X})$ to $\mathcal{C}_b(\mathcal{Z})$. Then, it follows from (73) that

$$|g_{n,f_1}(z, y) - g_{n,f_2}(z, y)| \leq \|A\|_\infty G(y) \cdot \|f_1 - f_2\|_\infty.$$

According to (van der Vaart and Wellner, 1996, § 2.7.4), this implies

$$N_{[]} (2\varepsilon \|A\|_\infty \|G\|_{L_2(P)}, \mathcal{G}_n, \|\cdot\|_{L_2(P)}) \leq N(\varepsilon, \mathcal{F}, \|\cdot\|_\infty) \quad (75)$$

where $N_{[]}(\cdot)$ denotes the bracketing number and $N(\cdot)$ denotes the covering number; see, e.g., (van der Vaart and Wellner, 1996, § 2.1.1). According to the assumptions on \mathcal{X} and k , it follows from (van der Vaart and Wellner, 1996, Theorem 2.7.1) that there is some $C_0 \in (0, \infty)$ such that, for every $\varepsilon > 0$, we have $\log N(\varepsilon, \mathcal{F}, \|\cdot\|_\infty) \leq C_0 \varepsilon^{-d/m}$; see, e.g., (Hable, 2012b, Eqn. (61)). Hence, it follows from (75) that

$$\log N_{[]}(\varepsilon, \mathcal{G}_n, \|\cdot\|_{L_2(P)}) \leq C_0 \cdot (2\|A\|_\infty \|G\|_{L_2(P)})^{\frac{d}{m}} \cdot \varepsilon^{-\frac{d}{m}} \quad \forall \varepsilon > 0.$$

Recall that $m > d/2$. Hence, the bracketing integral fulfills

$$J_{[]}(\delta_n, \mathcal{G}_n, \|\cdot\|_{L_2(P)}) = \int_0^{\delta_n} \sqrt{\log N_{[]}(\varepsilon, \mathcal{G}_n, \|\cdot\|_{L_2(P)})} d\varepsilon \xrightarrow{n \rightarrow \infty} 0 \quad (76)$$

for every sequence of real numbers $\delta_n \searrow 0$.

Finally, the definition of $g_{n,f}$, (70), and (52) imply $\lim_{n \rightarrow \infty} g_{n,f}(z, y) = 0$ for every $f \in \mathcal{F}$ and $(z, y) \in \mathcal{Z} \times \mathcal{Y}$ so that the dominated convergence theorem and (74) yield

$$\lim_{n \rightarrow \infty} \mathbb{E}[g_{n,f_1}(Z_i, Y_i)g_{n,f_2}(Z_i, Y_i)] - \mathbb{E}g_{n,f_1}(Z_i, Y_i)\mathbb{E}g_{n,f_2}(Z_i, Y_i) = 0 \quad (77)$$

for all $f_1, f_2 \in \mathcal{F}$. Summing up, because of (72), (74), (76), and (77), the assumptions in (van der Vaart, 1998, § 19.5) are fulfilled and it follows from (van der Vaart, 1998, Theorem 19.28) that

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n g_{n,f}(Z_i, Y_i) - \mathbb{E}_P g_{n,f} \right)_{f \in \mathcal{F}} \rightsquigarrow \mathbb{G}_P \quad \text{in } \ell_\infty(\mathcal{F})$$

where \mathbb{G}_P is a tight Gaussian process. That is, it only remains to prove $\mathbb{G}_P = 0$. This follows from considering finite marginals: Fix any $f_1, \dots, f_s \in \mathcal{F}$ and note that, due to (74) and (77), the (multivariate) Lindeberg-Feller central limit theorem (e.g., van der Vaart, 1998, Prop. 2.27) for the random vectors

$$(g_{n,f_1}(Z_i, Y_i), \dots, g_{n,f_s}(Z_i, Y_i)), \quad i \in \{1, \dots, n\},$$

implies that $(\mathbb{G}_P(f_1), \dots, \mathbb{G}_P(f_s)) = 0$. \square

Proof of Prop. 6.4: According to Assumption (19) the set $H_0 := \{f \in H \mid \mathcal{R}_{A,P}(f) < \infty\}$ is non-empty. Since $B_f : [0, \infty) \rightarrow \mathbb{R}$, $\lambda \mapsto \mathcal{R}_{A,P,\lambda}(f)$ is continuous for every $f \in H_0$, the map $B : [0, \infty) \rightarrow \mathbb{R}$ defined by $B(\lambda) = \inf_{f \in H_0} B_f(\lambda)$, $\lambda \in [0, \infty)$, is upper semi-continuous. For every $n \in \mathbb{N}$, the function $f_{A,P,\lambda_n} \in H$ uniquely exists according to Theorem 2.4 and we have $B(\lambda_n) = \mathcal{R}_{A,P,\lambda_n}(f_{A,P,\lambda_n})$. Then, the statement follows from

$$\begin{aligned} \inf_{f \in H} \mathcal{R}_{A,P}(f) &\leq \liminf_{n \rightarrow \infty} \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) \leq \limsup_{n \rightarrow \infty} \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) \leq \\ &\leq \limsup_{n \rightarrow \infty} B(\lambda_n) \leq B(0) = \inf_{f \in H_0} \mathcal{R}_{A,P}(f) = \inf_{f \in H} \mathcal{R}_{A,P}(f) \end{aligned}$$

\square

Proof of Theorem 6.5: First of all, it is shown that

$$(f_n)_{n \in \mathbb{N}_0} \subset H, \quad f_n \xrightarrow{w} f_0 \quad \Rightarrow \quad \lim_{n \rightarrow \infty} \mathcal{R}_{A,P}(f_n) = \mathcal{R}_{A,P}(f_0). \quad (78)$$

According to Prop. 6.1, it follows from weak convergence of the sequence $(f_n)_{n \in \mathbb{N}_0} \subset H$ that $\lim_{n \rightarrow \infty} \|A(f_n) - A(f_0)\|_\infty = 0$. Hence, there is an $a \in (0, \infty)$ such that, for every $n \in \mathbb{N}_0$, we have $\|A(f_n)\|_\infty \leq a$. Then, (78) follows from

$$\begin{aligned} \left| \mathcal{R}_{A,P}(f_n) - \mathcal{R}_{A,P}(f_0) \right| &\stackrel{(12)}{\leq} \int (b'_0 + b'_1 a^p) \cdot |A(f_n) - A(f_0)| \, dP \\ &\leq \|A(f_n) - A(f_0)\|_\infty \cdot \int (b'_0 + b'_1 a^p) \, dP. \end{aligned}$$

Next, define

$$H_0 := \left\{ f^* \in H \mid \mathcal{R}_{A,P}(f^*) = \inf_{f \in H} \mathcal{R}_{A,P}(f) \right\} \stackrel{(53)}{\neq} \emptyset.$$

It follows from (78) that H_0 is closed in H . Furthermore, it follows from convexity of $t \mapsto L(z, y, t)$ that H_0 is a convex set. Hence, H_0 has a unique element $f_{A,P}$ of smallest norm; see, e.g., (Denkowski et al., 2003, Theorem 3.7.9). That is, we have shown unique existence of an element $f_{A,P} \in H$ which fulfills (54) and (55).

The following proof of (56) is similar to the proof of (Steinwart and Christmann, 2008, Theorem 5.17). Let us fix a sequence $(\lambda_n)_{n \in \mathbb{N}} \subset (0, \infty)$ such that $\lim_{n \rightarrow \infty} \lambda_n = 0$ and, first, prove

$$\|f_{A,P,\lambda_n}\|_H \leq \|f_{A,P}\|_H \quad \forall n \in \mathbb{N} \quad (79)$$

by contradiction. If (79) was not true, then $\|f_{A,P,\lambda_n}\|_H > \|f_{A,P}\|_H$ for some $n \in \mathbb{N}$ and it would follow from $f_{A,P} \in H_0$ that $\mathcal{R}_{A,P,\lambda_n}(f_{A,P,\lambda_n}) > \mathcal{R}_{A,P,\lambda_n}(f_{A,P})$, which is a contradiction to the definition of f_{A,P,λ_n} . That is, we have shown (79). Now, we prove (56) by contradiction. If f_{A,P,λ_n} does not converge to $f_{A,P}$, then there is an $\varepsilon > 0$ and a subsequence $(f_{A,P,\lambda_{n_\ell}})_{\ell \in \mathbb{N}}$ such that

$$\|f_{A,P,\lambda_{n_\ell}} - f_{A,P}\|_H > \varepsilon \quad \forall \ell \in \mathbb{N}. \quad (80)$$

Due to (79), the subsequence $(f_{A,P,\lambda_{n_\ell}})_{\ell \in \mathbb{N}}$ is bounded and, therefore, there is a further subsequence which converges weakly; see, e.g., (Dunford and Schwartz, 1958, Cor. IV.4.7). That is, we may assume without loss of generality that $(f_{A,P,\lambda_{n_\ell}})_{\ell \in \mathbb{N}}$ fulfills (80) and converges weakly to some $f_0 \in H$. Then, it follows from (78) that $\lim_{\ell \rightarrow \infty} \mathcal{R}_{A,P}(f_{A,P,\lambda_{n_\ell}}) = \mathcal{R}_{A,P}(f_0)$. Hence, Prop. 6.4 implies that $\mathcal{R}_{A,P}(f_0) = \inf_{f \in H} \mathcal{R}_{A,P}(f)$ and, due to (55),

we either have $\|f_0\|_H > \|f_{A,P}\|_H$ or $f_0 = f_{A,P}$. However, weak convergence implies

$$\|f_0\|_H \leq \liminf_{\ell \rightarrow \infty} \|f_{A,P,\lambda_{n_\ell}}\|_H \leq \limsup_{\ell \rightarrow \infty} \|f_{A,P,\lambda_{n_\ell}}\|_H \stackrel{(79)}{\leq} \|f_{A,P}\|_H$$

This proves $f_0 = f_{A,P}$ and $\lim_{\ell \rightarrow \infty} \|f_{A,P,\lambda_{n_\ell}}\|_H = \|f_{A,P}\|_H$. Finally, this latter convergence of the norms together with weak convergence of $f_{A,P,\lambda_{n_\ell}}$ to $f_0 = f_{A,P}$ implies $\lim_{\ell \rightarrow \infty} \|f_{A,P,\lambda_{n_\ell}} - f_{A,P}\|_H = 0$; see, e.g., (Conway, 1985, Exercise V.1.8). This is a contradiction to (80). \square

Proof of Theorem 3.1: Theorem 3.1 immediately follows from Theorem 6.3 and 6.5. \square

Proof of Theorem 3.2: Theorem 3.1 guarantees existence of $f_{A,P} \in H$ which fulfills (21) and (22). As in the case of ordinary regularized kernel methods (i.e. $A = \text{id}$),

$$\begin{aligned} 0 &\leq \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P}^* \leq \mathcal{R}_{A,P,\lambda_n}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P}^* \leq \\ &\leq \mathcal{R}_{A,P,\lambda_n}(f_{A,P}) - \mathcal{R}_{A,P}^* \stackrel{(21)}{=} \lambda_n \|f_{A,P}\|_H^2; \end{aligned}$$

see (Steinwart and Christmann, 2008, p. 182). Hence,

$$\left| \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P}^* \right| \leq \lambda_n \|f_{A,P}\|_H^2. \quad (81)$$

Fix $n_0 \in \mathbb{N}$ such that $a_n \geq 1$ for every $n \geq n_0$. For every $\varepsilon \in (0, 1)$, define

$$B_{\varepsilon,n} := \left\{ D_n \in (\mathcal{Z} \times \mathcal{Y})^n \mid a_n \|f_{A,D_n,\lambda_n} - f_{A,P,\lambda_n}\|_\infty < \varepsilon \right\} \quad \forall n \in \mathbb{N}$$

and $a := \sup_{n \in \mathbb{N}} \|f_{A,P,\lambda_n}\|_\infty + 1$. It follows from Theorem 6.5 and (9) that $0 < a < \infty$. Then, Assumption 2.2 implies that, for every $D_n \in B_{\varepsilon,n}$ and $n \geq n_0$,

$$\begin{aligned} a_n \left| \mathcal{R}_{A,P}(f_{A,D_n,\lambda_n}) - \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) \right| &\leq \\ &\leq \int b'_0 + b'_1 a^p dP \cdot a_n \|f_{A,D_n,\lambda_n} - f_{A,P,\lambda_n}\|_\infty \end{aligned}$$

Since $\lim_{n \rightarrow \infty} P^n(B_{\varepsilon,n}) = 1$ for every $\varepsilon \in (0, 1)$ according to Theorem 6.5 and (9), it follows that

$$\lim_{n \rightarrow \infty} P \left(\left\{ D_n \in (\mathcal{Z} \times \mathcal{Y})^n \mid a_n \left| \mathcal{R}_{A,P}(f_{A,D_n,\lambda_n}) - \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) \right| < \varepsilon \right\} \right) = 1$$

for every $\varepsilon > 0$. Then, (26) follows from (81) and (25).

Finally, an elementary calculation shows that (25) is fulfilled for $\lambda_n = \gamma n^{-1/(4+p)} \quad \forall n \in \mathbb{N}$ for a constant $\gamma \in (0, \infty)$. \square

Proof of Lemma 6.6: According to the model assumptions, there is a set $B \in \mathfrak{B}_{\mathcal{Z}}$ such that $P_{\mathcal{Z}}(B) = 1$ and, for every $z \in B$, the conditional distribution $P(\cdot|z)$ is equal to the distribution of $t_z^* + s(z)\varepsilon$ where $t_z^* = (Af_{A,P})(z)$ is the (unique) median of the conditional distribution $P(\cdot|z)$. Now fix any $z \in B$. Since t_z^* is the median,

$$P((-\infty, t_z^*] | z) \geq \frac{1}{2} \geq P((t_z^*, \infty) | z) \quad (82)$$

$$P([t_z^*, \infty) | z) \geq \frac{1}{2} \geq P((-\infty, t_z^*) | z) \quad (83)$$

and the model assumptions (28)–(32) imply

$$P\left((t_z^*, t_z^* + \delta) \mid z\right) \geq \frac{c_h}{\bar{c}_s} \cdot \delta \quad \forall \delta \in (0, \alpha) \quad (84)$$

$$P\left((t_z^* - \delta, t_z^*) \mid z\right) \geq \frac{c_h}{\bar{c}_s} \cdot \delta \quad \forall \delta \in (0, \alpha). \quad (85)$$

The rest of the proof is divided into different cases; first, we consider the case that $t > t_z^*$. Note that

$$|y - t| - |y - t_z^*| > 0 \quad \text{if } t_z^* < y < \frac{1}{2}(t_z^* + t) \quad (86)$$

$$|y - t| - |y - t_z^*| > t_z^* - t \quad \text{if } \frac{1}{2}(t_z^* + t) \leq y < t \quad (87)$$

Then, dividing the domain of integration by t_z^* , $\frac{1}{2}(t_z^* + t)$, and t into four parts yields

$$\begin{aligned} \int L(z, y, t) P(dy|z) - \int L(z, y, t_z^*) P(dy|z) &= \int |y - t| - |y - t_z^*| P(dy|z) \\ &\stackrel{(86),(87)}{\geq} (t - t_z^*) \cdot P\left((-\infty, t_z^*] \mid z\right) + 0 \cdot P\left((t_z^*, \frac{1}{2}(t_z^* + t)) \mid z\right) - \\ &\quad - (t - t_z^*) \cdot P\left((\frac{1}{2}(t_z^* + t), t) \mid z\right) - (t - t_z^*) \cdot P\left([t, \infty) \mid z\right) \\ &\stackrel{(82)}{\geq} (t - t_z^*) \cdot P\left((t_z^*, \frac{1}{2}(t_z^* + t)) \mid z\right) \stackrel{(84)}{\geq} \frac{c_h}{2\bar{c}_s} (t - t_z^*)^2. \end{aligned}$$

Next, we consider the case that $t < t_z^*$. Similarly as before, dividing the domain of integration by t , $\frac{1}{2}(t_z^* + t)$, and t_z^* into four parts yields

$$\begin{aligned} \int L(z, y, t) P(dy|z) - \int L(z, y, t_z^*) P(dy|z) &= \int |y - t| - |y - t_z^*| P(dy|z) \\ &\geq -(t_z^* - t) \cdot P\left((-\infty, t] \mid z\right) - (t_z^* - t) \cdot P\left((t, \frac{1}{2}(t_z^* + t)) \mid z\right) - \\ &\quad + 0 \cdot P\left((\frac{1}{2}(t_z^* + t), t_z^*) \mid z\right) + (t_z^* - t) \cdot P\left([t_z^*, \infty) \mid z\right) \\ &\stackrel{(83)}{\geq} (t_z^* - t) \cdot P\left((\frac{1}{2}(t_z^* + t), t_z^*) \mid z\right) \stackrel{(85)}{\geq} \frac{c_h}{2\bar{c}_s} (t_z^* - t)^2. \end{aligned}$$

Finally, the remaining case $t = t_z^*$ is trivial. \square

Proof of Theorem 3.3: First, note that P and L fulfill Assumption 2.2 for $p = 0$ and that the model assumptions (28)–(32) imply $\mathcal{R}_{A,P}(f_{A,P}) = \inf_{f \in H} \mathcal{R}_{A,P}(f)$.

According to (7) and Theorem 6.5, there is an $n_0 \in \mathbb{N}$ such that $\|Af_{A,P,\lambda_n} - Af_{A,P}\|_\infty \leq \|A\| \cdot \|f_{A,P,\lambda_n} - f_{A,P}\|_H < a_h c_s =: \alpha$ for every $n \geq n_0$. Hence, it follows from Lemma 6.6 and (81) in the proof of Theorem 3.2 that there is a constant $\tilde{c} \in (0, \infty)$ such that

$$\|Af_{A,P,\lambda_n} - Af_{A,P}\|_{L_2(P_{Z_0})} \leq \tilde{c} \lambda_n^{\frac{1}{2}} \quad \forall n \geq n_0. \quad (88)$$

According to Assumption (39),

$$g_0 := \sum_{j=1}^{\infty} \frac{\langle f_{A,P}, v_j \rangle_H}{\sigma_j} \cdot u_j \in L_2(P_{Z_0})$$

and it follows from (35) and the fact that $\{v_j \mid j \in \mathbb{N}\}$ is a complete orthonormal system of H that

$$f_{A,P} = A_0^* g_0. \quad (89)$$

An easy calculation shows

$$\|f_{A,P,\lambda_n} - f_{A,P}\|_H^2 - \|f_{A,P,\lambda_n}\|_H^2 = 2\langle f_{A,P} - f_{A,P,\lambda_n}, f_{A,P} \rangle_H - \|f_{A,P}\|_H^2. \quad (90)$$

According to the definitions, we have

$$\mathcal{R}_{A,P}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P}(f_{A,P}) \geq 0 \quad \forall n \in \mathbb{N} \quad (91)$$

$$\mathcal{R}_{A,P,\lambda_n}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P,\lambda_n}(f_{A,P}) \leq 0 \quad \forall n \in \mathbb{N}. \quad (92)$$

Recall that $\mathcal{R}_{A,P,\lambda_n}(f) = \mathcal{R}_{A,P}(f) + \lambda_n \|f\|_H^2$. Hence,

$$\begin{aligned} \lambda_n \|f_{A,P,\lambda_n} - f_{A,P}\|_H^2 &\leq \\ &\stackrel{(91)}{\leq} \mathcal{R}_{A,P}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P}(f_{A,P}) + \lambda_n \|f_{A,P,\lambda_n}\|_H^2 - \\ &\quad + \lambda_n \|f_{A,P,\lambda_n} - f_{A,P}\|_H^2 - \lambda_n \|f_{A,P,\lambda_n}\|_H^2 = \\ &\stackrel{(90)}{=} \mathcal{R}_{A,P,\lambda_n}(f_{A,P,\lambda_n}) - \mathcal{R}_{A,P,\lambda_n}(f_{A,P}) + 2\lambda_n \langle f_{A,P} - f_{A,P,\lambda_n}, f_{A,P} \rangle_H \leq \\ &\stackrel{(92),(89)}{\leq} 2\lambda_n \langle f_{A,P} - f_{A,P,\lambda_n}, A_0^* g_0 \rangle_H = 2\lambda_n \langle A_0 f_{A,P} - A_0 f_{A,P,\lambda_n}, g_0 \rangle_{L_2(P_{Z_0})} \\ &\leq 2\lambda_n \|A_0 f_{A,P,\lambda_n} - A_0 f_{A,P}\|_{L_2(P_{Z_0})} \cdot \|g_0\|_{L_2(P_{Z_0})}. \end{aligned}$$

Together with (88), this implies that there is a constant $c \in (0, \infty)$ such that

$$\|f_{A,P,\lambda_n} - f_{A,P}\|_H^2 \leq c \lambda_n^{\frac{1}{2}} \quad \forall n \geq n_0.$$

Hence, if $\lim_{n \rightarrow \infty} a_n \lambda_n^{\frac{1}{2}} = 0$, it follows from Theorem 6.3 that

$$a_n \|f_{A, \mathbf{D}_n, \lambda_n} - f_{A, P}\|_H^2 \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.}$$

Next, let Assumption (41) be fulfilled. Then,

$$g := \sum_{j=1}^{\infty} \frac{v_j(x)}{\sigma_j} \cdot u_j \in L_2(P_{\mathcal{Z}_0})$$

and

$$A_0^* g \stackrel{(35)}{=} \sum_{j=1}^{\infty} v_j(x) \cdot v_j \stackrel{(10)}{=} \sum_{j=1}^{\infty} \langle v_j, \Phi(x) \rangle_H \cdot v_j = \Phi(x)$$

where we have used in the last equality that $\{v_j \mid j \in \mathbb{N}\}$ is a complete orthonormal system of H . Hence,

$$\begin{aligned} f_{A, P, \lambda_n}(x) - f_{A, P}(x) &\stackrel{(10)}{=} \langle f_{A, P, \lambda_n} - f_{A, P}, \Phi(x) \rangle_H = \\ &= \langle f_{A, P, \lambda_n} - f_{A, P}, A_0^* g \rangle_H = \langle A_0 f_{A, P, \lambda_n} - A_0 f_{A, P}, g \rangle_{L_2(P_{\mathcal{Z}_0})} \end{aligned}$$

and, accordingly, it follows from (88) that there is a constant $c \in (0, \infty)$ such that

$$(f_{A, P, \lambda_n}(x) - f_{A, P}(x))^2 \leq c \lambda_n \quad \forall n \geq n_0.$$

Hence, if $\lim_{n \rightarrow \infty} a_n \lambda_n = 0$, it follows from Theorem 6.3 and the reproducing property (10) that

$$a_n (f_{A, \mathbf{D}_n, \lambda_n}(x) - f_{A, P}(x))^2 \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in probability.}$$

□

References

- F. Bauer and A. Munk. Optimal regularization for ill-posed problems in metric spaces. *Journal of Inverse and Ill-Posed Problems*, 15(2):137–148, 2007.
- N. Bissantz and M. Birke. Asymptotic normality and confidence intervals for inverse regression models with convolution-type operators. *Journal of Multivariate Analysis*, 100(10):2364–2375, 2009.
- N. Bissantz and H. Holzmann. Statistical inference for inverse problems. *Inverse Problems*, 24(3):034009, 17, 2008.
- N. Bissantz and H. Holzmann. Asymptotics for spectral regularization estimators in statistical inverse problems. *Computational Statistics*, 28:435–453, 2013.

- N. Bissantz, T. Hohage, A. Munk, and F. Ruymgaart. Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM Journal on Numerical Analysis*, 45(6):2610–2636, 2007.
- B. Caputo, K. Sim, F. Furesjo, and A. Smola. Appearance-based object recognition using SVMs: Which kernel should I use? In *Proceedings of NIPS workshop on Statistical methods for computational experiments in visual processing and computer vision*, Whistler, 2002.
- L. Cavalier. Nonparametric statistical inverse problems. *Inverse Problems*, 24:034004, 19, 2008.
- L. Cavalier. Inverse problems in statistics. In *Inverse problems and high-dimensional estimation*, volume 203 of *Lecture Notes in Statistics*, pages 3–96. Springer, Heidelberg, 2011.
- A. Christmann and I. Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 13:799–819, 2007.
- A. Christmann, A. Van Messem, and I. Steinwart. On consistency and robustness properties of support vector machines for heavy-tailed distributions. *Statistics and Its Interface*, 2:311–327, 2009.
- J. B. Conway. *A course in functional analysis*. Springer-Verlag, New York, 1985.
- K. De Brabanter, J. De Brabanter, J. A. K. Suykens, and B. De Moor. Approximate confidence and prediction intervals for least squares support vector regression. *IEEE Transactions on Neural Networks*, 22:110–120, 2011.
- Z. Denkowski, S. Migórski, and N. S. Papageorgiou. *An introduction to nonlinear analysis: Theory*. Kluwer Academic Publishers, Boston, 2003.
- N. Dunford and J. T. Schwartz. *Linear operators. I. General theory*. Wiley-Interscience Publishers, New York, 1958.
- P. P. B. Eggermont, V. N. LaRiccia, and M. Z. Nashed. Moment discretization for ill-posed problems with discrete weakly bounded noise. *International Journal on Geomathematics*, 3:155–178, 2012.
- H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*. Kluwer Academic Publishers Group, Dordrecht, 1996.
- R. Hable. Asymptotic confidence sets for general nonparametric regression and classification by regularized kernel methods. arXiv:1203.4354, 2012a. *Submitted*.

- R. Hable. Asymptotic normality of support vector machine variants and other regularized kernel methods. *Journal of Multivariate Analysis*, 106: 92–117, 2012b.
- R. Hable. Universal consistency of localized versions of regularized kernel methods. *Journal of Machine Learning Research*, 14:153–186, 2013.
- R. Hable and A. Christmann. On qualitative robustness of support vector machines. *Journal of Multivariate Analysis*, 102:993–1007, 2011.
- A. Karatzoglou, A. Smola, K. Hornik, and A. Zeileis. kernlab – an S4 package for kernel methods in R. *Journal of Statistical Software*, 11:1–20, 2004.
- J. Krebs. Support vector regression for the solution of linear integral equations. *Inverse Problems*, 27:065007, 2011.
- J. Krebs, A. K. Louis, and H. Wendland. Sobolev error estimates and a priori parameter selection for semi-discrete Tikhonov regularization. *Journal of Inverse and Ill-Posed Problems*, 17:845–869, 2009.
- R. Kress. *Linear integral equations*. Springer, New York, second edition, 1999.
- B. A. Mair and F. H. Ruymgaart. Statistical inverse estimation in Hilbert scales. *SIAM Journal on Applied Mathematics*, 56:1424–1444, 1996.
- P. Mathé and S. V. Pereverzev. Optimal discretization of inverse problems in Hilbert scales. Regularization and self-regularization of projection methods. *SIAM Journal on Numerical Analysis*, 38:1999–2021, 2001.
- D. W. Nychka and D. D. Cox. Convergence rates for regularized solutions of integral equations from discrete noisy data. *The Annals of Statistics*, 17:556–572, 1989.
- F. O’Sullivan. A statistical perspective on ill-posed inverse problems. *Statistical Science*, 1(4):502–527, 1986.
- R. T. Rockafellar. *Convex analysis*. Princeton University Press, Princeton, N.J., 1970.
- R. T. Rockafellar. Integral functionals, normal integrands and measurable selections. In *Nonlinear operators and the calculus of variations*, pages 157–207. Springer, Berlin, 1976.
- B. Schölkopf and A. J. Smola. *Learning with kernels*. MIT Press, Cambridge, 2002.

- I. Steinwart and A. Christmann. *Support vector machines*. Springer, New York, 2008.
- A. van der Vaart. *Asymptotic statistics*. Cambridge University Press, Cambridge, 1998.
- A. van der Vaart and J. Wellner. *Weak convergence and empirical processes. With applications to statistics*. Springer, New York, 1996.
- A. C. M. van Rooij and F. H. Ruymgaart. Asymptotic minimax rates for abstract linear estimators. *Journal of Statistical Planning and Inference*, 53(3):389–402, 1996.
- V. N. Vapnik. *Statistical learning theory*. John Wiley & Sons, New York, 1998.
- G. Wahba. Practical approximate solutions to linear operator equations when the data are noisy. *SIAM Journal on Numerical Analysis*, 14:651–667, 1977.
- G. Wahba. Ill posed problems: Numerical and statistical methods for mildly, moderately, and severely ill posed problems with noisy data. Technical Report 595, University of Wisconsin-Madison, Department of Statistics, 1980.
- H. Wendland. *Scattered data approximation*. Cambridge University Press, Cambridge, 2005.